MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LIBRARIES

M.I.T. LIBRARIES - DEWEY

HB31
.M415
no.
94-
32

Dewey

working paper
department
of economics

EXPLAINING INVESTMENT DYNAMICS IN U.S. MANUFACTURING: A GENERALIZED (S,s) APPROACH

Ricardo J. Caballero

94-32

Sept. 1994

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

EXPLAINING INVESTMENT DYNAMICS IN U.S. MANUFACTURING: A GENERALIZED (S,s) APPROACH

Ricardo J. Caballero

94-32 Sept. 1994

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

OCT 21 1994

LIBRARIES

Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach

BY RICARDO J. CABALLERO AND EDUARDO M.R.A. ENGEL[1]

First version: September, 1993

This version: September, 1994

In this paper we derive a model of aggregate investment that builds from the lumpy microeconomic behavior of firms facing stochastic fixed adjustment costs. Instead of the standard (S, s) bands, firms' optimal adjustment policies are probabilistic, with a probability of adjusting (*adjustment hazard*) that grows smoothly with firms' disequilibria. Depending upon the specification of the distribution of fixed adjustment costs, the adjustment hazards approach encompasses models ranging from the very non-linear (S, s) model to the linear partial adjustment model. Except for the latter extreme, the processes for aggregate investment obtained from adding up the actions of firms subject to aggregate and idiosyncratic shocks, is highly non-linear. Estimating the aggregate model by maximum likelihood, we find clear evidence supporting non-linear models over linear ones for postwar sectoral U.S. manufacturing equipment and structures investment. For a given sequence of aggregate shocks, the nonlinear model estimated generates brisker expansions and — to a lesser extent — sharper contractions than its linear counterpart. These features fit well the observed positive skewness and large kurtosis of U.S. manufacturing sectoral investment/capital ratios.

---

[1] MIT and NBER; Universidad de Chile and NBER. We are grateful to Olivier Blanchard, Whitney Newey, James Stock, and seminar participants at Brown, CEPR-Champoussin, Econometric Society Meeting in Caracas, Harvard, IMPA (Rio), NBER EFCC, Princeton, Rochester, U. Católica (Ing. Sist.), U. de Chile (Centro de Economía Aplicada) and Yale for their comments. Ricardo Caballero thanks the National Science and Sloan foundations, and Eduardo Engel FONDECYT (Chile), Grant 92/901, for financial support.

Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

http://www.archive.org/details/explaininginvest00caba

1 INTRODUCTION

Characterizing the dynamic behavior of aggregate investment has not been easy. Variables that should be significant, such as cost of capital and q, seldom are; variables that should not appear in investment equations, such as cash flows and income, often do. Moreover, estimates are typically unstable, differing across samples, estimation procedures, and researchers.[2] Not everything is so dismal, however, for recent work has improved the state of affairs, mainly by refining the measurement and instrumenting of cost of capital,[3] and by incorporating credit constraints into the analysis.[4]

Common to the new and old literatures is a limited treatment of dynamics. This is somewhat surprising, for there seems to be consensus on adjustment costs playing an important role in determining aggregate investment dynamics. Indeed, investment functions — rather than stock demands — are interesting and well defined objects only if frictions in stock adjustments are significant. In spite of this consensus, adjustment costs and investment dynamics have been mostly treated as a nuisance to deal with when trying to estimate the effect of cost of capital (or other relevant variables) on capital accumulation. A clear example of this tradition can be found in Hall and Jorgenson's (1967) seminal paper, where an elegant frictionless demand for capital theory is transformed into an investment theory by means of a few ad hoc lags, which then turn out to play a crucial role in estimation. The same criticism applies to most q models. The elegance of q-theory derives from two features: First, barring measurement issues, q is a sufficient statistic for all information about future demand and productivity conditions affecting the firm. And second, the mapping from q to investment depends exclusively on the adjustment cost function. Although the first of these features is regularly stressed, the latter is often disregarded and researchers limit their study to linear or log-linear relations between q and investment. Of course, by now we know how to justify these specifications, both the lags and the linear q-models; the empirical workhorse is the quadratic adjustment costs model which, besides "validating" these specifications, delivers simple smooth and linear dynamics.[5]

Quadratic adjustment costs may be a useful approximation when investment is just a part of a more general model, but they are difficult to justify when studying investment in isolation. It is even more difficult to give any structural interpretation to the resulting estimates, for their implications are at odds with the basic microeconomic investment facts. At the plant level, leaving aside minor upgrades and repairs, investment is intermittent and lumpy rather than smooth. This is starkly documented in Doms and Dunne (1993). They use the Longitudinal Research Datafile to study the investment behavior of 12,000 continuing (and large) U.S. manufacturing establishments for the seventeen year period from 1972-1988, and find that (i) more then half of the establishments exhibit capital growth between 40 and 60 percent in a single year, and (ii) between 25 and 40 percent of a plant's gross investment over the seventeen year period is concentrated in a single year.[6]

---

[2] See Chirinko (1994) for a survey of the empirical investment literature. See Clark (1979), Bernanke (1983), and Blanchard, Rhee, Summers (1993) for "horse races."

[3] See e.g. Auerbach and Hassett (1992), Clark (1993).

[4] See e.g. Fazzari, Hubbard and Petersen (1988), Hubbard and Kashyap (1992), Bernanke, Gertler and Gilchrist (1993).

[5] See Rothschild (1971) for a compelling discussion of the relevance of more general adjustment costs, in particular non-convexities, for investment dynamics.

In this paper we focus our attention entirely on the dynamic aspects of aggregate investment. In doing so we impose the constraint that the theory must be consistent with the basic lumpy and intermittent nature of microeconomic data. These constraints yield several methodological advantages. Among these, they facilitate a meaningful structural interpretation of our findings, therefore inheriting the standard advantages of structural parameters.[7] The results in this paper, however, go beyond methodological considerations and show that modeling inaction at the microeconomic level makes a difference at the aggregate level. We characterize the aggregate nonlinearities implied by the model and study their role in shaping aggregate investment dynamics.

At the microeconomic level, there has been extensive development of models of lumpy and intermittent adjustment (the (S,s) literature).[8] Here we extend these models so the adjustment trigger barriers vary randomly across firms and for a firm over time. This modification is a first step toward introducing the realistic and empirically important feature that units do not always wait for the same stock disequilibrium to adjust, and that adjustments are not always of the same size across firms and for the same firm over time.

Recently, there have also been developments of empirical models of aggregate dynamics with heterogeneous microeconomic units adjusting intermittently.[9] Econometric implementation of these models, however, has required observing (or estimating separately in a first stage) a measure of the aggregate driving force. In the current context, this amounts to constructing a cost of capital measure. But undoubtedly many of the problems of the empirical investment literature are due to the difficulties of constructing a proper measure of the cost of capital, a variable that despite the current efforts is likely to be plagued by simultaneity and omitted variables problems.[10] In this paper we implement a nonlinear time series procedure that does not require the first stage; it only requires information on the investment series itself and on the generating process of the driving force (but not its realization). Somewhat analogously with the standard procedure of estimating convex adjustment costs parameters from the first (or higher) order serial correlation of investment, we learn about more complex and realistic lumpy adjustment cost functions from the structure of aggregate investment lags and their changes over time.

---

[6] Since plants entry is excluded from their sample, these statistics are likely to represent lower bounds on the degree of lumpiness in plants' investment patterns.

[7] See Bertola and Caballero (1990) for a similar motivation. Another important advantage of paying attention to microeconomic aspects of adjustment when modeling aggregate adjustment, is that this facilitates integrating microeconomic and macroeconomic data for aggregate purposes, see Caballero and Engel (1993b), Caballero, Engel and Haltiwanger (1993), and Eberly (1994).

[8] See Harrison, Sellke and Taylor (1983) for a technical discussion of impulse control problems. For a good survey of the economics literature — although with an emphasis on models where investment is infrequent but not lumpy — see Dixit and Pindyck (1994). A model more closely related to a special case of ours is Grossman and Laroque's (1990) model of consumer durable purchases.

We estimate nonlinear dynamic panel data models for two-digit U.S. manufacturing investment/capital ratios for the period 1948-1992. We find clear and widespread evidence in favor of our generalized (S,s) model over simple linear models. Perhaps the most revealing result occurs when a generalized (S,s) model with parameters constrained to be equal across all sectors (3 parameters), outperforms a linear model with *unrestricted* (across sectors) AR(2) processes (42 parameters). Our structural interpretation of these non-linearities indicates that the fraction of firms' capital subject to fixed costs is large, and that these costs are also large. Although important for both, these features are more pronounced for structures than equipment.

One of the main mechanisms by which aggregate dynamics generated by (S,s) type models differs from their linear counterpart, is that the number of active firms changes over the cycle — a point emphasized by Bar-Ilan and Blinder (1992). Doms and Dunne's (1993) confirm the importance of this mechanism; they show that the number of plants going through their primary investment spikes, rather than the average size of these spikes, tracks closely aggregate manufacturing investment over time. For a given sequence of aggregate shocks, the nonlinear model estimated in this paper generates brisker expansions and — to a lesser extent — sharper contractions than its linear counterpart. These features fit well the observed positive skewness and large kurtosis of U.S. sectoral investment/capital ratios.

---

[9] Blinder (1981), Bar-Ilan and Blinder (1992) and Lam (1992) look at data on inventories (the first one) and consumer durables (the other two) under the organizing principles of (S,s) models. Bertola and Caballero (1990) and Caballero (1993) provide a structural empirical framework and estimate (S,s) models for consumer durable goods. Bertola and Caballero (1994) implement empirically an irreversible investment model where microeconomic investment is intermittent but not lumpy. Caballero and Engel (1992a, 1993b, 1993c) estimate aggregate models of employment and price adjustments when microeconomic units follow more general (probabilistic) microeconomic adjustment rules but, contrary to the current paper, they do not derive these rules from a microeconomic optimization problem.

[10] For a discussion of some of these biases see Shapiro (1986) and Blanchard's (1986) comment on that paper. Also see Cummins, Hassett and Hubbard (1994) for an argument relying on measurement error and noise to focus estimation on periods where shocks are known to be large. They find that estimated adjustment costs are much smaller in these periods than in non-tax credit reform periods. Part of their finding may be due to the nonlinearities we describe in this paper.

The next section presents the basic model. Section 3 describes the econometric method. Our main empirical results are presented and discussed in section 4. Section 5 extends the basic model to allow for flexible as well as fixed capital. Conclusions are presented in Section 6.

2 THE BASIC MODEL

2.1 Overview

We model a sector composed of a large but fixed number of monopolistically competitive firms. Each firm faces an isoelastic demand for its differentiated product, which is produced with a Cobb-Douglas constant returns technology on labor and capital. Both demand and technology are affected by multiplicative shocks described by a joint geometric random walk process. These shocks have firm specific and sectoral (aggregate) components which we specify later. We work in discrete time.

The sector faces infinitely elastic supplies of labor and capital. We choose the price of the latter as numeraire and let the wage (in terms of capital) follow a geometric random walk process, possibly correlated with demand and technology shocks. Firms can adjust their labor input at will but suffer a loss when resizing their stock of capital. We assume that this loss is an increasing function of the firm's scale of operation;[11] it can be interpreted either as an index of the degree of specificity of firms' capital, or as a secondary market imperfection if machines or structures are replaced, or as a reorganization cost associated with putting new capital to work.

For a given scale of operation, the extent of the loss due to adjustment varies over time as firms may, for example, find better or worse matches or uses for their old machines, or may face reorganizations of different degree of difficulty. For simplicity, but at the cost of realism, we model the proportional loss factor as a random variable independent across firms and time, and we assume its realization is known by the firm when it decides whether or not to resize its current stock of capital.

---

[11] The precise meaning of "scale of operation" depends on the specific interpretation of the adjustment cost.

As in standard (S,s) models, the resulting microeconomic policy is one of inaction interspersed with periods of large investment or disinvestment. As in standard search models, at each point in time the firm decides whether to "accept" the currently offered adjustment cost (proportional loss factor) or to postpone adjustment and draw a new adjustment cost next period. The interaction between these two mechanisms implies that, more realistically than in standard (S,s) models, the size of adjustments varies both across firms and over time for the same firm. During a given time period, firms with identical shortages or excesses of capital act differently; over time, the same firm reacts differently to similar disequilibria in its stock of capital.

Intuitively, the largest adjustment cost for which a firm does not adjust its stock of capital decreases with the extent of its capital stock imbalance. If the distribution of adjustment costs is non-degenerate, this implies that the probability that a firm adjusts given the firm's disequilibrium — a concept we describe as the firm's *adjustment hazard* — increases smoothly and monotonically with the firm's disequilibrium in its stock of capital.[12]

Since the adjustment cost factors are independent across firms, and we assume that the number of firms is large, the *adjustment hazard* described above characterizes the sectoral investment at each point in time. Given firms' capital imbalances at the beginning of a period, the fraction of units resizing their stock of capital is determined by the adjustment hazard. Sectoral investment is the sum of the products of the adjustment hazard and the size of the investment undertaken by those firms that decide to adjust. Equivalently, it is the sum of the expected investment by firms, conditional on their capital stock imbalances.

Except for the degenerate case where the adjustment hazard is constant (partial adjustment model), sectoral investment depends critically on the number of firms at each position in the space of capital imbalances, thereby motivating our focus on the *cross sectional density of disequilibria.* The dynamics of sectoral investment are then determined by the evolution of the cross sectional density of disequilibria. The path of this density is driven by the interaction of sectoral, firm specific, and adjustment cost shocks with the history of shocks and actions contained in previous cross sectional densities of disequilibria.

---

[12] This should be contrasted with standard (S,s) models, where the probability of a firm adjusting jumps from zero to one at the trigger points, and also with standard linear partial adjustment models, where this probability is independent of the size of the firm's disequilibrium. Later we argue that these polar extremes correspond to two degenerate cases: the standard (S,s) model to the case where the distribution of adjustment costs collapses to a mass point, the partial adjustment model to the case where the distribution of adjustment costs degenerates in the opposite direction (infinite variance keeping the location parameter fixed).

2.2 The firm

When the firm is not investing, its flow of profits is:

$$\Pi(K,\theta) = K^{\beta}\theta - (r+\delta)K, \tag{1}$$

where K is the firm's stock of capital; θ is a shock to the profit function which combines demand, productivity and wage shocks;[13] r and δ are the discount and depreciation rates; and $\beta \equiv (\alpha(\eta-1))/(1+\alpha(\eta-1))$ with α the elasticity of production with respect to capital and η the elasticity of demand faced by the monopolist. As is standard, $\eta > 1$ for the monopolist solution to be interior; which implies that $\beta < 1$, since we assume that $\alpha < 1$.

It is useful to replace θ in the profit function by a variable with more economic content. We do this by defining the frictionless stock of capital of the firm, K^*, as the solution of the maximization of (1) with respect to capital, so that:

$$\theta = \xi K^{*(1-\beta)},$$

where $\xi \equiv (r+\delta)/\beta$. Substituting this expression into (1), and defining the disequilibrium variable

$$z \equiv \ln K/K^*,$$

allows us to rewrite the profit function as:

$$\Pi(z,K^*) \equiv \pi(z)K^* = \xi\left(e^{\beta z} - \beta e^{z}\right)K^*. \tag{2}$$

Figure 1 illustrates, and equation (2) implicitly defines, profits *per unit of frictionless capital*, $\pi(z)$.

At times of adjustment, the firm incurs in an adjustment cost proportional to foregone profits due to reorganization.[14] Assuming this cost corresponds to the opportunity cost

---

[13] For convenience, we have written the profit function net of flow payment on capital, $(r+\delta)K$. Since there are neither borrowing constraints nor bankruptcy options, the solution to the firm's problem is unchanged by replacing flow payments for a lump sum payment at the time of purchase.

[14] See e.g. Cooper and Haltiwanger (1993).

GAUSS Thu Nov 4 15:43:39 1993

Figure 1

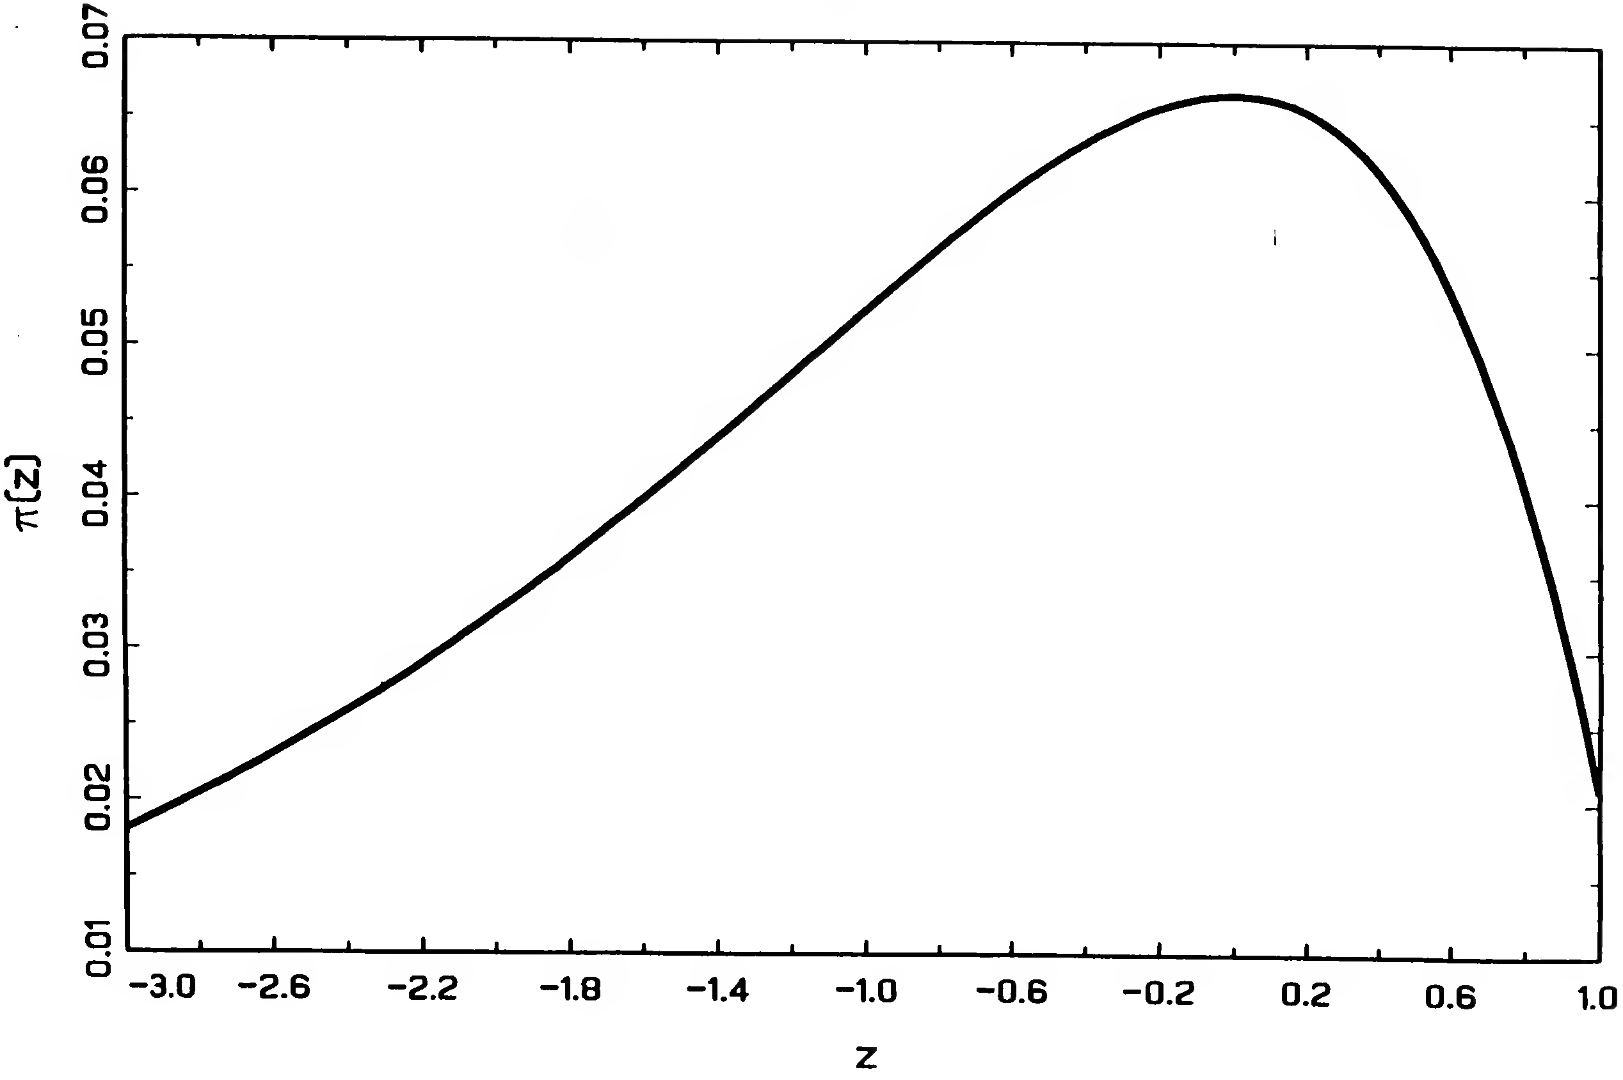

of installed capital during reorganization, a derivation similar to the one that led to (2) implies that:

$$\text{Adjustment Cost} = K^{\beta}\theta = \omega\xi e^{\beta z^{-}} K^{*},$$

where ω is the realization of a positive random variable and z^{-} denotes the capital imbalance immediately before adjustment.[15] Realizations of ω are generated by a common distribution function, $G(\omega)$, and are independent across firms and over time.[16]

Given the increasing returns nature of the adjustment cost technology, the optimal policy is obviously not one of continuous and small investments but rather one of periods of inaction followed by occasional lumpy investment. Therefore, the firm's problem can be characterized in terms of two regimes: action and inaction. Finding a solution to the firm's problem means finding the function separating these two regions in (ω, z)-space, and characterizing the actions taken by the firm when crossing from the inactive to the active region. We turn to this next.

The value of a firm with disequilibrium z, frictionless stock of capital K^{*}, and (current) adjustment cost parameter ω, $V^{*}(z, K^{*}, \omega)$, is the maximum of the value of the firm if it does not adjust, $V(z, K^{*})$, and the value if it does adjust, $V(c, K^{*}) - \omega\xi e^{\beta z} K^{*}$; where c is the optimally determined return point (see below). In short:

$$V^{*}(z_t, K_t^{*}, \omega_t) = \max\{V(z_t, K_t^{*}),\, V(c, K_t^{*}) - \omega_t \xi e^{\beta z_t} K_t^{*}\}. \tag{3}$$

The evolution of the value of a firm that does not adjust in the current period is described by:

$$V(z_t, K_t^{*}) = \pi(z_t) K_t^{*} + (1+r)^{-1} \mathrm{E}_t[V^{*}(z_{t+1}, K_{t+1}^{*}, \omega_{t+1})]. \tag{4}$$

Since the profit and adjustment costs functions are homogeneous of degree one with respect to K^{*}, given z, so are the value functions $V(z, K^{*})$ and $V^{*}(z, K^{*}, \omega)$. This allows us to reduce the number of state variables by restating the problem in terms of the value per unit of frictionless capital. Let $v(z) \equiv V(z, K^{*})/K^{*}$ and $v^{*}(z, \omega) \equiv V^{*}(z, K^{*}, \omega)/K^{*}$. Dividing both sides of equations (3) and (4) by K^* and noting that

---

[15] Of course many other specifications of lumpy adjustment costs are possible. For example, these could be proportional to the frictionless or new capital rather than the old one. The former specification would simplify the problem since it would make the upgrading and downgrading decisions symmetric; the latter specification, on the other hand, would complicate the problem somewhat since the return points from upgrading and downgrading would be different. None of these modifications would change anything fundamental in what follows, however.

[16] We take this as a first step toward a more realistic formulation where at the individual level adjustment costs exhibit some persistence and, at any point in time, the distribution of adjustment costs depends on aggregate conditions.

$$\frac{K^*_{t+1}}{K^*_t} = (1-\delta)e^{-\Delta z_{t+1}},$$

yields

$$v^*(z_t,\omega_t) = \max\{v(z_t)\,,\, v(c) - \omega_t \xi e^{\beta z_t}\}, \tag{5}$$

$$v(z_t) = \pi(z_t) + \psi \mathrm{E}_t\left[v^*(z_{t+1},\omega_{t+1})e^{-\Delta z_{t+1}}\right], \tag{6}$$

with $\psi \equiv (1-\delta)/(1+r)$. Before going further, we use figure 2 to illustrate the basic setup developed up to now. This figure shows how $v(z)$, $v(c) - \omega\xi e^{\beta z}$, and $v^*(z,\omega)$ determine the trigger points, given a particular realization of the adjustment cost. The solid line illustrates the value of a firm that does not adjust in the current period. The concave dashed line represents the value of a firm that decides to adjust, given a realization of ω. The maximum between both lines describes $v^*(z,\omega)$, and the inaction range — for a given ω — corresponds to the interval between the intersection of the two lines.

It follows directly from maximization of the value of a firm that decides to adjust, $v(c) - \omega\xi e^{\beta z}$, with respect to the return point c, that the maximum of $v(z)$ and $v^*(z,\omega)$ is obtained at $z = c$ and that this return point is independent of the initial disequilibrium.

We let $\Omega(z)$ denote the largest adjustment cost factor for which the firm finds it advantageous to adjust given a capital imbalance z. From the value matching condition that equates the two terms on the right hand side of equation (5), it follows that:

$$\Omega(z) = \xi^{-1}e^{-\beta z}(v(c) - v(z)). \tag{7}$$

It should be apparent from the formulae and figure 2 that once $v(z)$ is known, the solution to the firm's problem is easily obtained. The Bellman equation for this function can be obtained by substituting equation (5) into (6):

$$v(z_t) = \pi(z_t) + \psi \mathrm{E}_t\left[e^{-\Delta z_{t+1}} \max\{v(z_{t+1})\,,\, v(c) - \omega_{t+1}\xi e^{\beta z_{t+1}}\}\right], \tag{8}$$

where E_t denotes the expectation, based on information available at time t, with respect to Δz_{t+1} and ω_{t+1}. Using the definition of $\Omega(z)$ and calculating the expectation with respect

GAUSS Thu Nov 4 16:06:51 1993

Figure 2

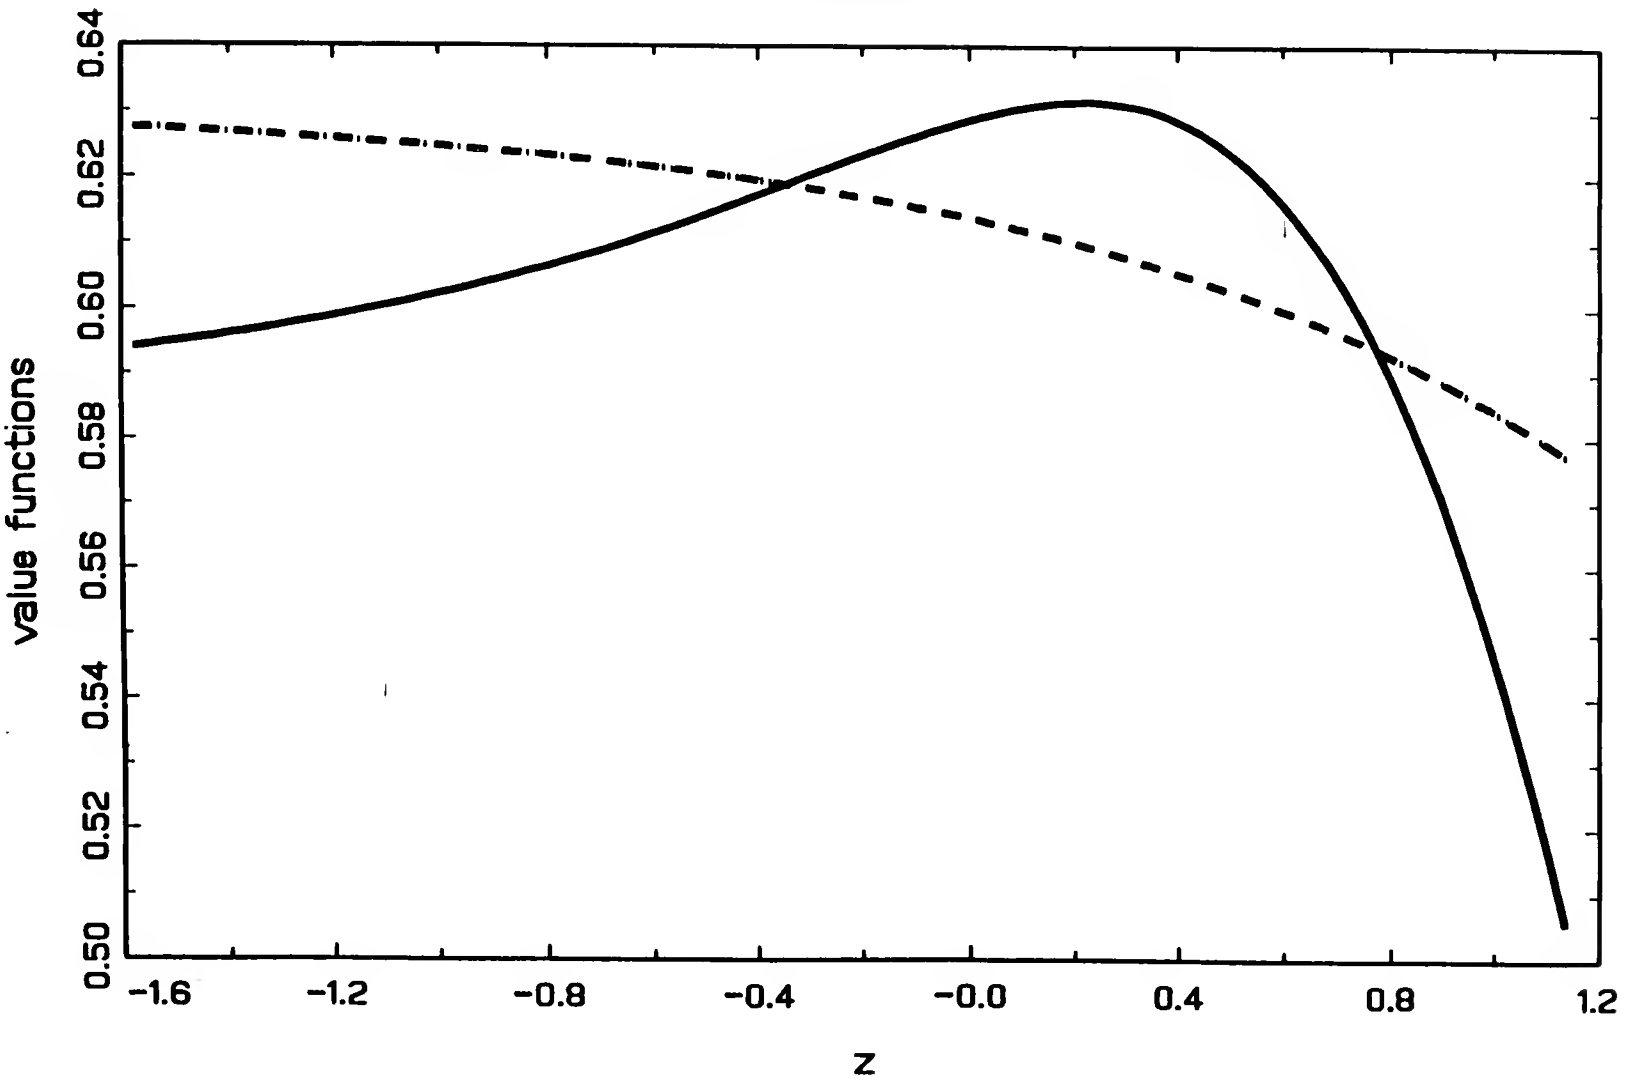

to ω_{t+1} leads to:

$$(9) \qquad v(z) = \pi(z) + \psi \mathrm{E}\left[e^{-\Delta z} \left\{ v(z+\Delta z) + \xi e^{\beta(z+\Delta z)} \int_0^{\Omega(z+\Delta z)} G(\omega)\, d\omega \right\} \right],$$

where $G(\omega)$ denotes the cumulative distribution function of the adjustment cost factor, and the expectation is taken over Δz. The first term inside the expectation can be interpreted as tomorrow's value of the firm for those realizations of the adjustment cost factor and shocks to z that lead to no investment, weighted by the probability of no adjustment for the corresponding Δz; the second term corresponds to the probability weighted value of those situations where the firm adjusts its stock of capital.

The value function when not adjusting, $v(z)$, can be found by solving the functional equation that results from replacing $\Omega(z)$ from (7) in (9), and using the first order condition $v'(c) = 0$. Alternatively, we can replace $v(z)$ and $v(c)$ from (9) in (7), in order to directly find a functional equation for the policy function $\Omega(z)$. After a few algebraic steps, the latter strategy yields:

$$\begin{aligned} \Omega(z) \;=\; & \psi \mathrm{E}\left[e^{-(1-\beta)\Delta z} \left\{ \int_0^{\Omega(z+\Delta z)} (1-G(\omega))\, d\omega - e^{\beta(c-z)} \int_0^{\Omega(c+\Delta z)} (1-G(\omega)) d\omega \right\} \right] + \\ & \xi^{-1} e^{-\beta z} \{\pi(c) - \pi(z)\}. \end{aligned}$$

The optimal policy $\Omega(z)$ is the solution to this functional equation subject to the first order condition with respect to c:

$$\Omega'(c) = 0,$$

which is obtained from differentiating (7) with respect to z, evaluating the resulting expression at $z = c$, using the fact that $\Omega(c) = 0$, and imposing the first order condition $v'(c) = 0$.

Figure 3a illustrates the function $\Omega(z)$ for an example where the distribution of adjustment costs is exponential: $dG(\omega) = (1/\lambda)\, e^{-\omega/\lambda}\, d\omega$. As follows from equation (7), if the firm's disequilibrium is close enough to $z = c$, it will adjust for arbitrarily small adjustment costs. From then on, $\Omega(z)$ increases with respect to $|z - c|$, more sharply to the left (capital shortages) than to the right (excess of capital). This asymmetry has several sources: First, since depreciation is strictly positive (and possibly net productivity and demand growth is on average positive), shortages are less likely to reverse by themselves than excesses, thus there is less value in delaying upgrading than downgrading. Second, by making the

Figure 3a

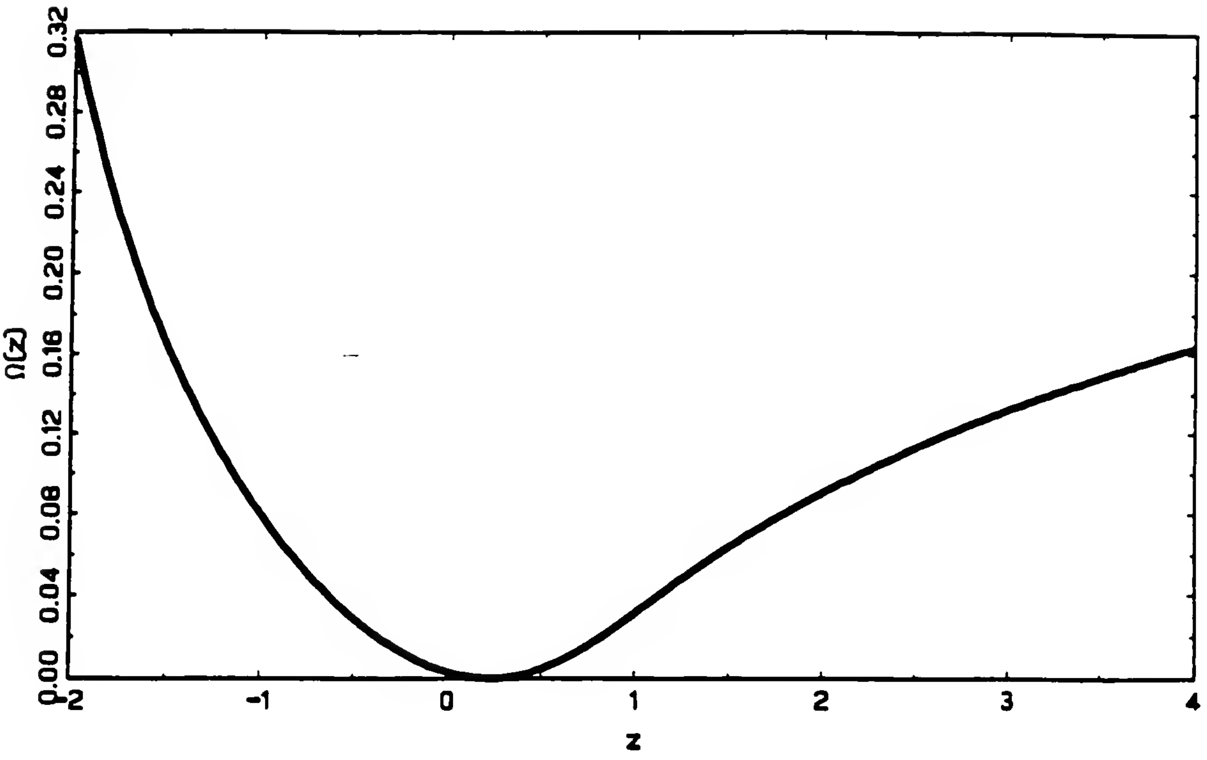

Figure 3b

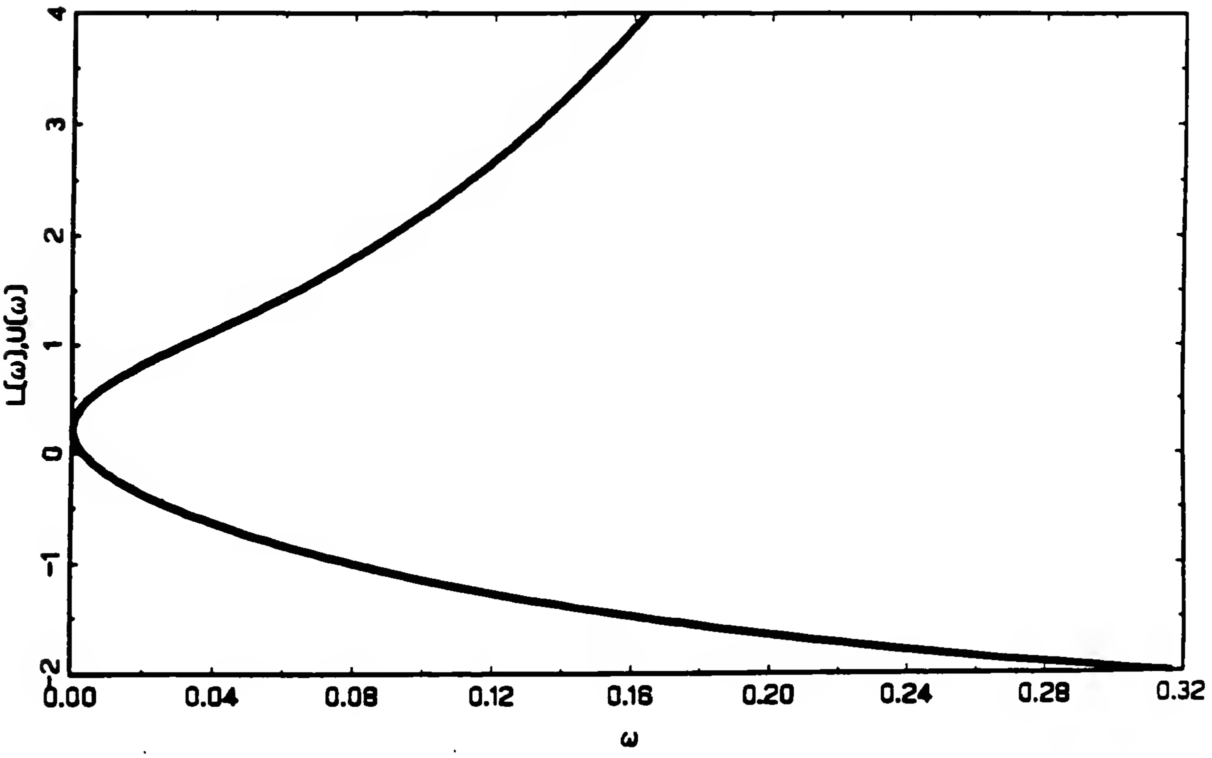

adjustment cost an increasing function of *old* capital (for a given disequilibrium) we have made the total adjustment cost asymmetric: for given K^* and ω, it is more costly to adjust from the right than from the left of c because old capital is larger in the former than in the latter case.[17] Third, the profit function is asymmetric; it decreases faster to the right of $z = c$ than to the left of its maximum (see figure 1). In our example, the first two sources of asymmetry outweigh the last one.

Figure 3b depicts the inverse of the function $\Omega(z)$. We label the segments of the curve below and above c, $L(\omega)$ and $U(\omega)$, respectively. These functions correspond to the maximum shortage and excess of capital tolerated by the firm for any given realization of the adjustment cost factor ω. That is, for any fixed ω, they describe a standard (L, c, U) policy.[18] The area enclosed by the two curves corresponds to the combinations of disequilibria and adjustment cost factors for which the firm chooses not to adjust.

The shape and location of the function $\Omega(z)$ and its inverse, $(L(\omega), c, U(\omega))$, depend on the entire distribution of adjustment cost factors, $G(\omega)$. A given realization of the adjustment cost factor will not generate the same inaction range for different distribution functions $G(\omega)$. In particular, a low value of ω is more likely to lead to action when it comes from a distribution of adjustment cost factors with a high rather than a low average value. We discuss the relation between $G(\omega)$ and microeconomic adjustment in more detail in the next section, when we characterize aggregate investment and its connection with the underlying distribution of adjustment cost factors.

2.3 Sectoral investment

2.3.1 The adjustment hazard and cross sectional distribution

Let $x \equiv z - c$ denote a firm's imbalance with respect to its target point rather than its frictionless stock of capital. Also let K_t^A, I_t^A, $K_t(x)$ and $I_t(x)$ denote the aggregate (sectoral) stock of capital and gross investment, and the stock of capital and gross investment held by firms with disequilibrium x at time t (before adjustment).

Those firms with deviation x whose current adjustment cost is small enough to make adjusting profitable (i.e. for which $\omega < \Omega(x + c)$) adjust. Since adjustment costs are i.i.d.

---

[17] This "irreversible investment" feature of the adjustment cost function seems realistic, lending further support to our choice of adjustment cost specification.

[18] An (L, c, U) policy corresponds to a two-sided (S, s) model. The notation L, U and c stands for lower bound, upper bound, and "center," respectively. See Harrison et al. (1983).

and the number of firms is large, the fraction of firms with deviation x that adjusts is approximately equal to:

$$\Lambda(x) = G(\Omega(x + c)), \tag{10}$$

where $G(\omega)$ denotes the cumulative distribution function for the adjustment cost factor ω. For example, if $G(\omega)$ is a *Gamma* distribution with mean $p\phi$ and variance $p\phi^2$, the adjustment hazard is:

$$\Lambda(x) = \frac{1}{\phi^p \Gamma(p)} \int_0^{\Omega(x+c)} \omega^{p-1} e^{-\omega/\phi}\, d\omega.$$

A firm with capital stock x has a probability $\Lambda(x)$ of adjusting its stock of capital and, if it does so, its contribution to aggregate investment is equal to $(e^{-x} - 1)K_t(x)$. Thus aggregate investment is given by:

$$I_t^A = \int (e^{-x} - 1)\Lambda(x)\overline{K_t(x)} f^*(x,t)dx,$$

where $\overline{K_t(x)}$ denotes the average stock of capital of firms with imbalance x and $\tilde{f}(x,t)$ the cross sectional density of disequilibria just before adjustments take place.

Assuming that the average capital stock of firms with disequilibrium x is approximately independent of x, we have that:[19]

$$\frac{I_t^A}{K_t^A} = \int (e^{-x} - 1)\Lambda(x)\tilde{f}(x,t)\, dx. \tag{11}$$

Equation (11) is the fundamental aggregate investment/capital ratio expression; it shows how, conditional on the cross-section of disequilibria, current investment is determined by the adjustment hazard function. Figure 4a shows the adjustment hazard function for the three different gamma distributions of adjustment cost factors depicted in figure 4b. These distributions have the same median but different variances: the solid line corresponds to $p = 1$ (exponential distribution) and $\phi = 0.1$, the long dashes correspond to a high variance distribution, while the short dashes describes a low variance distribution. Figure 4a shows that as the variance increases the hazard shifts from an (S, s) policy, to a smoothly increasing hazard, and eventually approaches a constant probability of adjustment/partial adjustment model.[20] The connection between these results and the height of the density of

---

[19] This independence assumption can be motivated by the fact that x is a stationary variable while K is not (because K^* is not). Provided firms have been in the industry for a long time, knowing a firm's x conveys no information about its level of capital.

[20] Strictly speaking, we obtain the partial adjustment model only if $e^{-x} - 1 \approx -x$ for most relevant x's.

Figure 4a: Adjustment hazards

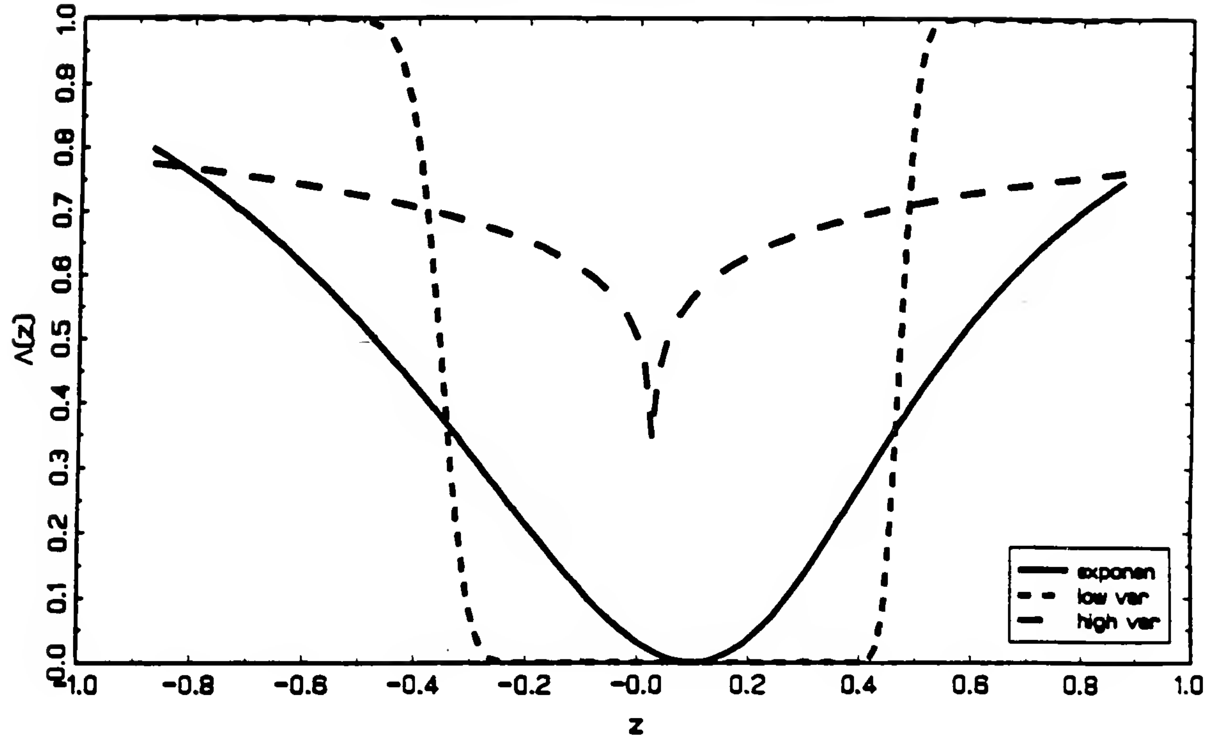

Figure 4b: Adjustment costs

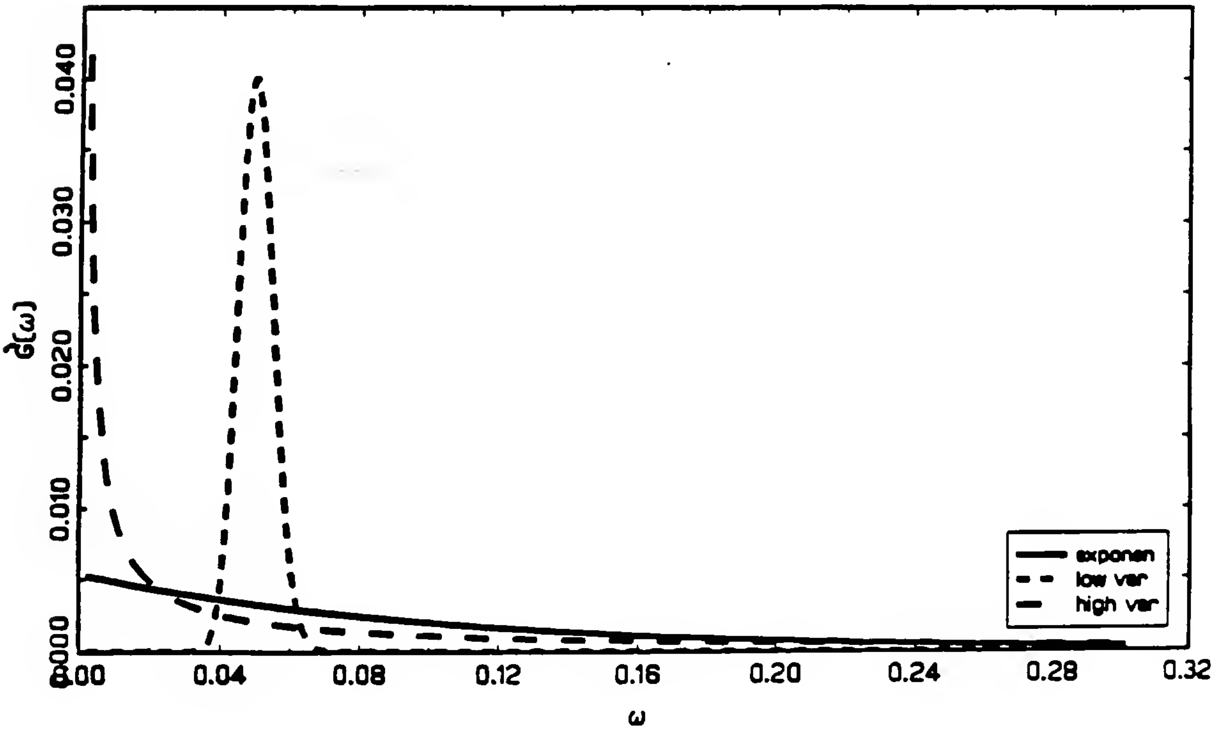

the adjustment cost around zero is apparent. If the density around zero is small, there is a range where the firm almost never adjusts since adjustment costs are almost never small enough to justify it; the standard (S, s) — or (L, c, U) — case is an extreme version of this. Conversely, when the density around zero is large — as is frequently the case when the variance grows keeping the median constant — the decision of adjustment is largely motivated by the adjustment cost draw rather than by the firm's disequilibrium. The limit of this, when adjustment becomes independent of the disequilibrium, corresponds to the partial adjustment model.

2.3.2 Sectoral equilibrium and cross sectional dynamics

Shocks to wages, demand, and productivity drive the dynamics of frictionless capital. We decompose these shocks into sectoral shocks, v_t, and firm specific (idiosyncratic) shocks, ϵ_t:

$$K_t^* = K_{t-1}^* e^{v_t + \epsilon_t},$$

which implies that when the firm does not adjust, the disequilibrium measure x evolves according to:

$$\Delta x_t = -v_t + \ln(1-\delta) - \epsilon_t \approx -(\delta + v_t) - \epsilon_t.$$

We assume that these shocks are exogenous to the firm and the sector.[21]

Between two consecutive periods, the cross section distribution of disequilibria changes as a result of firms' adjustments, depreciation, sectoral shocks, and idiosyncratic shocks. Since we are working in discrete time, it is important to describe the timing of events within each period. We denote the cross section density at the end of period $t-1$ by $f(x, t-1)$. Depreciation and the aggregate shock corresponding to period t follow, resulting in the density $\tilde{f}(x, t)$. Adjustments, as determined by the hazard function $\Lambda(x)$ come next, and period t concludes with the idiosyncratic shocks. The final density is $f(x, t)$, and the cycle starts again. Recalling that a positive shock leads to a decrease in x, we can summarize this chain of events as follows:

---

[21] The former assumption needs no explanation while the latter requires additional restrictions in the underlying monopolistic competitive model in order to eliminate the sectoral price from the demand curve faced by the firm. This requires precise offsetting of the income and substitution effects induced by sectoral price changes. An example of such a model can be found in Caballero and Engel (1993a). In any event, we view the lack of a full (or more general) equilibrium setup as a limitation of the current approach.

$$\begin{aligned}
(12) \quad \tilde{f}(x,t) &= f(x+\delta+v_t, t-1), \\
(13) \quad f(x,t) &= \left[\int \Lambda(y)\tilde{f}(y,t)dy\right] g_\epsilon(-x) + \int (1-\Lambda(x+\epsilon))\tilde{f}(x+\epsilon,t)g_\epsilon(-\epsilon)\,d\epsilon,
\end{aligned}$$

where $g_\epsilon(\epsilon)$ is the probability density for the idiosyncratic shocks. The integro-difference equation describing the evolution of the cross sectional distribution from one period to the next follows directly from equations (12) and (13):

$$\begin{aligned}
(14) \qquad f(x,t) &= \left[\int \Lambda(y)f(y+\delta+v_t,t-1)dy\right] g_\epsilon(-x) \\
&\quad + \int (1-\Lambda(x+\epsilon))f(x+\epsilon+\delta+v_t,t-1)g_\epsilon(-\epsilon)\,d\epsilon.
\end{aligned}$$

Combining equations (11), (12) and (14) we can fully characterize aggregate investment, for any given sequence of aggregate shocks $\{v_t\}$ and initial cross section distribution $f(x,0)$. We turn to estimation issues next.

3 ECONOMETRICS

3.1 Overview

The econometric problem consists of estimating the parameters of: (a) firms' profit functions, (b) the initial distribution of disequilibria, (c) the distribution of adjustment costs, (d) the distribution of idiosyncratic shocks, and (e) the process generating aggregate shocks. The data available are sectoral (aggregate) investment data.

We assume that the stochastic nature of the problem (the "error term") is due to aggregate shocks. If these shocks were known exactly, the fit for the "true" parameters would be perfect.[22] The aggregate shocks may either be totally unobserved, in which case we have to posit some underlying time series process captured by a subvector of Θ, the vector of all parameters in the model; or they may be related to observables via an auxiliary economic model or long run relation (both of which include residual error terms). Even though the approach we describe next can be generalized in a straightforward manner, and we do so later in the paper, to simplify the exposition here we consider the case where the aggregate shock is unobserved and i.i.d.

---

[22] Conditional on also knowing the initial cross sectional distribution. We assume there is no measurement error in y_t, which was the only source of error allowed for in previous analyses (see footnote 9). Unfortunately, integrating both approaches is computationally too burdensome.

The observed investment data may now be viewed as a highly nonlinear transformation of the unobserved aggregate shocks. Denoting I_t^A/K_t^A by y_t, we have that:

$$y_t = y_t(v_1, \ldots, v_t; \Theta); \qquad t = 1, \ldots, T, \tag{15}$$

where T denotes the number of time periods considered.

Whatever estimation criterion we use,[23] it will require that, for a given set of parameters, we compute the set of aggregate shocks that correspond to the observed sectoral investment data (and some functions of these shocks). This amounts to finding the inverse functions of those described in (15):[24]

$$v_t = v_t(v_1, \ldots, v_{t-1}, y_t; \Theta); \qquad t = 1, \ldots, T. \tag{16}$$

The effect of previous shocks on v_t is summarized in the cross-section density after adjustment in period $t-1$ takes place, so that:

$$v_t = v_t(f(\cdot, t-1), y_t); \qquad t = 1, \ldots, T. \tag{17}$$

The recursive nature of this problem is clear. The aggregate shock at time t cannot be calculated without knowing all the preceding shocks. Yet the effect of preceding shocks is summarized in the cross sectional density of firms' deviations *before* the current shock, as can be seen from replacing $\tilde{f}(x,t)$ by $f(x,t-1)$ in (11) and using (12):

$$y_t = \int (e^{-x+\delta+v_t} - 1)\Lambda(x - \delta - v_t) f(x, t-1)\, dx. \tag{18}$$

It follows that the computational problem of calculating the v_t defined in (16) corresponds to the following recursive problem:[25]

- Find v_1 by solving (18) with $f(x,t-1)$ replaced by the initial density (determined by the subvector of Θ that characterizes this distribution).

---

[23] Without loss of generality as far as the methodology we descibe next is concerned, we consider maximum likelihood estimation.

[24] And making sure that this inverse is uniquely determined, which we do shortly.

[25] To ensure uniqueness when solving (18), we show that the right hand side of this equation is an increasing function of v for any density $f(x,t-1)$. This is equivalent to showing that $\phi(u)\Lambda(-u)$ is increasing in u, where in our case $\phi(u) = e^u - 1$, but more generally, $\phi(u)$ could be any increasing function with $\phi(0) = 0$. A straightforward calculation shows that a sufficient condition for $\phi(u)\Lambda(-u)$ to be increasing in u is that $\Lambda(u)$ be decreasing ($\Lambda'(u) \leq 0$) for negative values of u and increasing ($\Lambda'(u) \geq 0$) for positive values of u. All the adjustment hazards used in the estimation section satisfy this condition.

- Determine $f(x, 1)$ by using equation (14).
- Find v_2 by solving (18) with $f(x, t-1)$ replaced by $f(x, 1)$.
- Determine $f(x, 2)$ by using equation (14).
- Find v_3 by solving (18) with $f(x, 2)$ in the place of $f(x, t-1)$.

$$\vdots$$

This problem entails a highly nonlinear but structured relation between the data and aggregate shocks. At each point in time, given the distribution of disequilibria, there is a nonlinear relation between the current aggregate shock, v_t, and the observation, y_t. The current shock also enters nonlinearly in the equation determining the cross sectional distribution of disequilibria relevant for next period's investment decisions; thus, establishing a nonlinear relation between observations and lagged aggregate shocks.

For given parameters, the sequence of aggregate shocks and the derivatives of aggregate investment with respect to these shocks (the inverse Jacobian), which are the inputs to the likelihood function, can in principle be obtained from the combination of a routine to solve the nonlinear static equation relating current shocks to current observations, and a stochastic Markov chain designed to track down the evolution of the cross sectional distribution. In practice, however, this approach is computationally burdensome for any Markov chain grid fine enough to prevent a host of numerical problems from arising.

Obviously, computational considerations are particularly important when the time series considered are long or when multiple time series are analyzed simultaneously, as we do later in the paper. In order to cope with this problem, in the estimation section we first consider a family of hazard functions which is general enough to test the basic implications of our model and, at the same time, simple enough so that an accurate approximation method that reduces the estimation time by several orders of magnitude can be designed. Next we consider a more general family of hazard functions, which provides good approximations for adjustment hazards obtained via value iteration for a wide range of parameter values. This family is used to estimate structural parameters.

3.2 A simple family of hazards

Estimating a structural model is computationaly burdensome. Thus, it seems sensible to start by estimating faster approximations which may shed light on the potential success of the proper structural model.

The algorithm for computing the aggregate shocks described above is considerably faster if we work with a (sufficiently rich) family of adjustment hazards, characterized by a small number of parameters, for which the cross-section densities remain within the family. What makes finding such a family of densities difficult is that the cross section distributions undergo three very different transformations from one period to the next, two of which are non-trivial. At the beginning of the new period, the aggregate shock shifts the cross section; this is followed by firms' adjustments (the hazard shock); and the period concludes with the convolution of the distribution resulting from the previous transformations with the distribution of idiosyncratic shocks.

The family of cross sectional densities we use is the family of mixture of normal densities. It is easy to see that this family is closed under aggregate shocks and idiosyncratic shocks, as long as we assume that the latter are also normal. To ensure that the cross section remains within this family after the hazard shock, we assume that the adjustment hazard is an inverted normal:

$$\Lambda(x) = 1 - e^{-\lambda_0 - \lambda_2 x^2}, \tag{19}$$

where $\lambda_0 \geq 0$ and $\lambda_2 \geq 0$.[26]

We will show that tracking the evolution of the cross sectional distribution reduces to keeping track of the mean, variance, and weights of a finite number of Normals. These parameters can be obtained from a simple recursive structure. Furthermore, the static nonlinear equation relating the contemporaneous aggregate shock to the current observation and cross sectional density is also much easier to solve, since evaluating the right hand side of (18) for different values of v_t is considerably faster.

In what follows we describe the basic issues and procedure with a simple univariate example. We focus on a maximum likelihood procedure and assume that the aggregate shock is unobservable and distributed Normal with mean μ_v and standard deviation σ_v. None of these assumptions is important for the substantive issues discussed below. With minor modifications, the procedure can be adapted to more robust (but less efficient) estimation procedures (e.g. GMM), to other distributional assumptions about aggregate shocks, to partial observability of the aggregate shock, and to a multivariate context (which we do later). We assume that idiosyncratic shocks are normal with zero mean and variance σ_ϵ^2 and, for simplicity, proceed conditionally on the initial cross section density; at the end of this section we discuss how we deal with initial conditions.

---

[26] It is straightforward to add a linear term, $-\lambda_1 x$, to the exponent.

The basic algorithm is conceptually simple. A change of variable calculation relating $(v_1, \ldots, v_T)$ to $(y_1, \ldots, y_T)$ shows that minus the log-likelihood is equal to:

$$-l(\Theta|y_1, \ldots, y_T) = \text{const} + \sum_t \ln\left|\frac{\partial y_t}{\partial v_t}\right| + \frac{T}{2}\ln\left(\sum_t \frac{(v_t-\mu)^2}{T}\right), \tag{20}$$

which is to be minimized over Θ, with y_t and v_t defined in (15) and (16). The first term captures the rich dynamics of the model under consideration. Conditional on the cross section density before the aggregate shock, current investment is a non-linear, increasing function of the current aggregate shock. The exact shape of this non-linearity depends on the cross sectional density at the time of the shock. Thus $\partial y_t/\partial v_t$ not only varies with v_t, but also over time. This should be contrasted with a partial adjustment model or an ARIMA model, where $\partial y_t/\partial v_t$ does not depend on the current shock and remains constant over time.

Next we show that assuming that the cross section density is a convex combination of normal densities, and the hazard an inverted normal, simplifies the calculation of the likelihood.

Let us begin by considering the simplest possible case, where the cross section density at time t is normal with mean μ and variance σ^2. A simple but tedious calculation, based upon (18), shows that finding the current shock (for the parameter vector Θ under consideration) is equivalent to solving for v_t in:

$$y_t = e^{c(v_t)} - 1 - \frac{\tau}{\sigma}e^{-d(v_t)+\frac{\tau^2}{\sigma^2}c(v_t)}, \tag{21}$$

where

$$\tau^2 = \frac{\sigma^2}{1+2\lambda_2\sigma^2}, \tag{22}$$

$$c(v) = v-\mu+\frac{1}{2}\sigma^2, \tag{23}$$

$$d(v) = \lambda_0 + \frac{\tau^2}{\sigma^2}\lambda_2(v-\mu)^2. \tag{24}$$

The corresponding derivative needed in the first term of the log-likelihood is equal to:

$$\frac{\partial y_t}{\partial v_t} = \mathrm{e}^{c(v_t)} + \frac{\tau^3}{\sigma^3}\mathrm{e}^{-d(v_t)} - 2\lambda_2(v_t - \mu)\frac{\tau^2}{\sigma^2}\left[\mathrm{e}^{c(v_t)} - y_t - 1\right].$$

Once v_t is found by solving (21), equation (14) can be used to show that the cross section density after the t-th period's aggregate, hazard and idiosyncratic shocks will be a convex combination of two normal densities, one of them with mean $\eta = (\mu - v_t)/(1 + 2\lambda_2\sigma^2)$ and variance $\tau^2 + \sigma_\epsilon^2$, and the other with zero mean and variance σ_ϵ^2. The former corresponds to those firms that did not adjust, the latter to those that adjusted their capital stock. The fraction of firms in the group that does not adjust is equal to:

$$\kappa = \frac{\tau}{\sigma}\exp\left(-\lambda_0 + \frac{(\mu - v)^2}{2\sigma^2}\left[\frac{\tau^2}{\sigma^2} - 1\right]\right).$$

In the more general case, where $f(x,t) = \sum_k \alpha_k f_k(x,t)$ is a convex combination of normal densities, v_t is obtained by solving an equation analogous to (21) with a linear combination of terms like the one on the right hand side of that equation:

$$y_t = \sum_k \alpha_k \int (\delta + v_t - x)\Lambda(x - \delta - v_t) f_k(x,t)dx.$$

The evolution of the cross-section density can be tracked by considering a linear combination of expressions like those considered above. For example, in the case where $f(x,t)$ is equal to the convex combination of N_t normal densities and $\Lambda(x)$ is given by (19), $f(x,t+1)$ is a convex combination of $N_t + 1$ normal densities. Each of these corresponds to a specific cohort, grouped according to the last time they adjusted. The "older" distributions are more spread out than the "younger" ones and have lost mass monotonically due to the adjustment of their members. In order to keep the number of normal densities considered manageable, we reduce the number of densities tracked down at each point in time by one, by merging the two older cohorts.

Finally, we need to determine the initial cross sectional density. The ergodic density is, in a precise sense, the best guess for this density (see Caballero and Engel, 1992b). We disregard the first T_0 observations when calculating the likelihood, thus allowing for the possible distortions introduced by this approximation to wash away.

We approximate the ergodic density by a convex combination of normal densities as follows. Let f_0^e denote the density of idiosyncratic shocks,[27] and f_1^e the density that results from those firms that do not adjust after f_0^e is subject to one set of aggregate, hazard, and idiosyncratic shocks.[28] Let f_2^e denote the density that results after applying an aggregate, hazard, and idiosyncratic shock to f_1^e, and so on. Standard Markov chain arguments show that the ergodic density has the form:

$$f_E(x) = \sum_{k\geq 0} \phi_k f_k^e(x),$$

where the ϕ_k's correspond to the stationary distribution of the Markov chain describing the number of periods since a firm last adjusted. Since the ϕ_k's decrease at a geometric rate, a relatively small number of densities usually provides a good approximation. Figure 5 depicts the corresponding weighted densities and the resulting ergodic density. The largest density corresponds to the density of idiosyncratic shocks, the density that follows to the left is the density of those that do not adjust after a sequence of aggregate, idiosyncratic and hazard shocks, and so on.

3.3 A general family of hazards

Experimentation with a variety of distributions of adjustment costs shows that the family of continuous, piecewise inverted normal hazards approximates well the hazard functions obtained via value iteration. Three pieces suffice for most practical purposes, with the middle piece corresponding to a hazard that is identically equal to zero. A representative member of this four-parameter family of adjustment hazards is of the form:

$$\Lambda(x) = \begin{cases} 1 - e^{-\lambda^-(x-x^-)^2} & \text{if } x < x^-, \\ 0 & \text{if } x^- \leq x \leq x^+, \\ 1 - e^{-\lambda^+(x-x^+)^2} & \text{if } x > x^+. \end{cases} \tag{25}$$

In this case we track the evolution of the cross section density and solve the non-linear equations that provide the v_t's by working with a flexible discretization in x-space. We describe this discretization procedure next.

Assume the cross section at the beginning of period $t-1$ is a collection of mass points: (x_i, w_i), $i = 1, \ldots, n$, where x_i denotes where the i-th mass point is located and w_i the

---

[27] Here idiosyncratic shocks are as faced by an individual firm. Thus the corresponding variance is equal to $\sigma_\epsilon^2 + \sigma_v^2$. See Caballero and Engel (1992b) for details.

[28] The size of the aggregate shock is equal to the mean of this process.

FIGURE 5

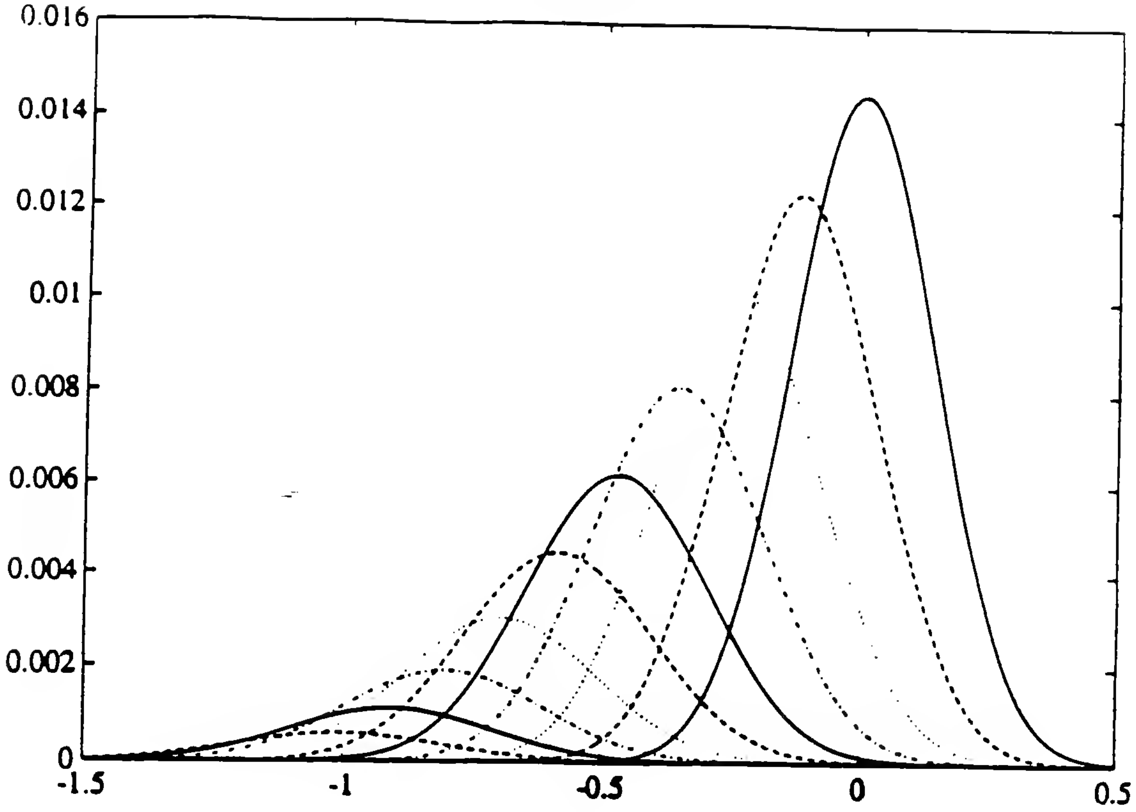

0.016
0.014
0.012
0.01
0.008
0.006
0.004
0.002
0
-1.5
-1
-0.5
0
0.5

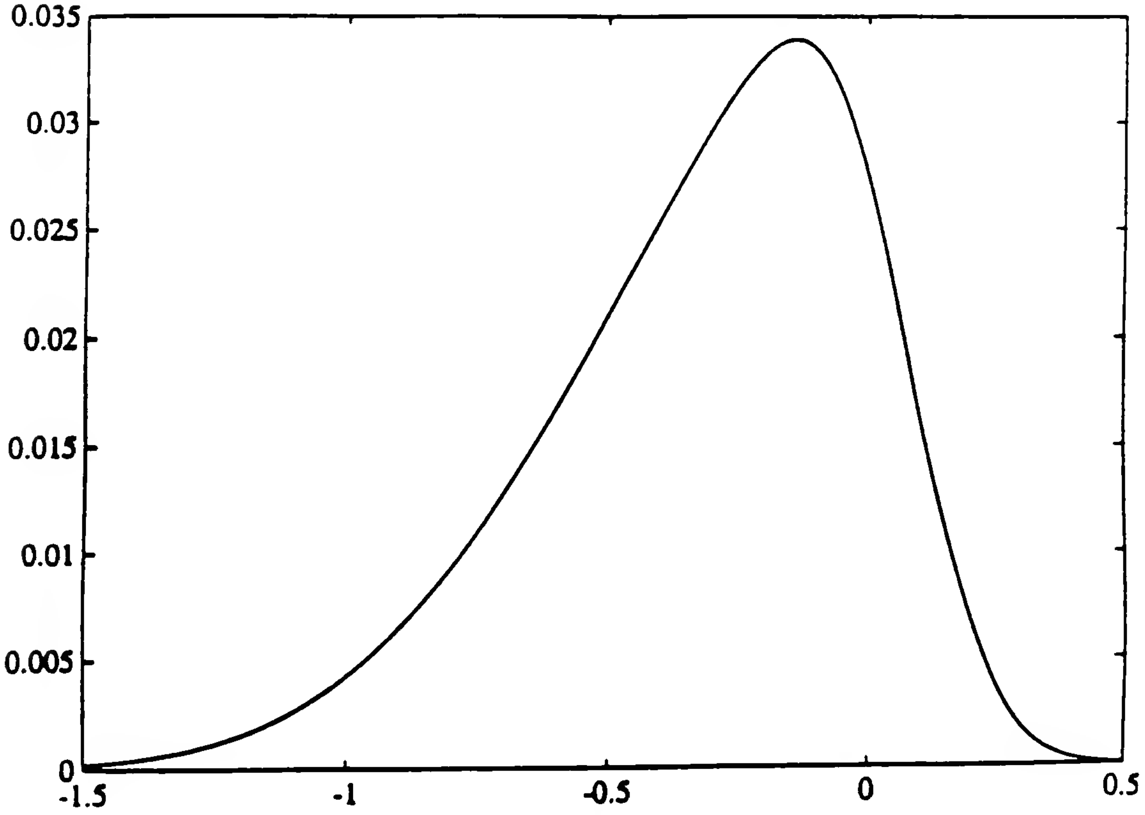

0.035
0.03
0.025
0.02
0.015
0.01
0.005
0
-1.5
-1
-0.5
0
0.5

corresponding weight ($\sum_i w_i = 1$). The aggregate shock v_t is obtained by solving the discrete counterpart of equation (18). Having found v_t, the evolution of the cross section proceeds as follows: After the aggregate shock hits, the resulting collection of mass points is characterized by $(x_i - v_t, w_i)$. Next follows the hazard shock, which increases the number of mass points by one, since a mass point at zero arises. The idiosyncratic shock follows, leading to a continuous distribution. This distribution is discretized by evaluating the cross section density on a grid whose width and location are determined by the current mean and variance.[29] The grid has n points so that the cycle described above begins again.

Even though this methodology is considerably slower than the one described in the preceding subsection, it provides a practical way of approximating a rich family of hazard functions, thereby making structural estimation possible. Given a distribution of adjustment costs, the corresponding "exact" hazard is obtained via value iteration. We then approximate this hazard by the "closest" hazard within the family described by (25). We use this hazard in the iterative procedure described above, which calculates the sequence of aggregate shocks corresponding to particular parameter values.

4 EMPIRICAL EVIDENCE: U.S. INVESTMENT

4.1 Overview and Data

Our data are constructed from annual gross investment and capital series for 21 two-digit manufacturing industries from 1947 to 1992.[30] All series are in 1987 dollars, and the stock of capital correspond to the series used by the Bureau of Labor Statistics for their productivity studies.[31] Since capital stocks are end-of-year, our measures of the investment/capital ratio used in estimation start in 1948. We report separate results for equipment and structures panels; each has 945 observations.

We begin this section by estimating simple adjustment hazard functions for sectoral (aggregate) U.S. manufacturing equipment and structures investment, without imposing the tight theoretical restrictions of the microeconomic model developed in section 2. From this semi-structural approach, we learn that there is clear trace in aggregate data of microeconomic adjustment hazard functions that are increasing rather than constant with respect to firms' disequilibria. We then proceed to estimate structural models, and find that the average adjustment cost (in terms of annual revenues) is about 16 percent for equipment and 59 percent for structures, and that the standard deviation of these costs is 6 and 8 percent, respectively. Our estimate of the standard deviation of idiosyncratic (within each sector) shocks ranges from 5 to 12 percent and is quite stable across investment categories and estimation methods (structural and semi-structural).

---

[29] In the estimation section we use 33 equally spaced points between $\mu - 4\sigma$ and $\mu + 4\sigma$, where μ and σ denote the mean and standard deviation of the density being approximated.

[30] We have 21 rather than 20 sector because Motor Vehicles is separated from Transportation equipment.

[31] This is one of the three capital stock series reported by the Bureau of Economic Activity.

4.2 Results: Semi-structural hazard

Using the panels described above, we estimate the following "inverted-normal" hazard:

$$\Lambda(x) = 1 - e^{-\lambda_0 - \lambda_2 x^2}, \tag{26}$$

which has the potential to capture the increasing nature of the adjustment hazards depicted in section 3, and converges to the partial adjustment model as λ_2 approaches zero (together with the approximation $(e^{-x} - 1) \approx -x$). We restrict the hazard to be the same across sectors within each panel, but allow for different means, variances and first order correlations of sectoral shocks. We also allow for a general variance-covariance matrix across sectors' innovations.

For each category (equipment and structures), we solve at each point in time 21 static nonlinear equations (one for each sector) to find the (sectoral) aggregate shocks; with these on hand, we use 21 dynamic equations to generate new sectoral cross sectional distributions. After repeating these steps for all observations, we compute the serial correlation coefficients, the corresponding sectoral innovations, the variance-covariance matrix of sectoral innovations, and the derivatives of aggregate investment with respect to the current shock. Using these results, we maximize the multivariate likelihood over λ_0, λ_2 and σ_ϵ.[32]

We constrain the initial distributions of disequilibria to be the sectoral ergodic distributions corresponding to the parameters being estimated. In order to reduce the effect of the approximation introduced by this criterion, we exclude the first 3 shocks of each sector from the joint likelihood. Finally, at each point in time we keep track of 20 distributions for each sector and merge the distributions of those that have spent more than 20 years without adjusting.[33]

---

[32] I.e., we concentrate out the variance-covariance matrix and serial correlation coefficients.

[33] Remember that each distribution represents a cohort of the firms that have not adjusted for the last 1, 2, 3,..., 20 (or more) periods.

Table 1 displays the results. The two major subdivisions correspond to equipment and structures. Within each of these, the first column shows the estimates for a partial adjustment model, the second column reports the results for the inverted-normal hazard in equation (26), and the third column shows the estimates obtained with unconstrained (across sectors) second-order autoregressions for every series.[34]

For equipment, the first column shows a coefficient λ_0 which is significantly different from zero and implies a partial adjustment coefficient of 0.44.[35,36] The mean first order serial correlation of the sectoral shocks is 0.19, with a range between −0.03 and 0.47. The second column contains the main results of this section. These clearly indicate the relevance of considering non-constant hazards: The likelihood function exhibits a substantial improvement when λ_2 is allowed to be positive,[37] while the linear term, λ_0, becomes indistinguishable from zero (as implied by our model, see equations (7) and (10)).[38] At the same time, the average serial correlation of sectoral shocks is (on average) reduced (it ranges from −0.13 to 0.23). Comparing these results with those in the third column, reveals that the likelihood obtained with *unconstrained* second-order autoregressions leads to a likelihood that is substantially *lower* than that of the constrained (across sectors) non-linear model.

If we take the setup we used in the section describing the microeconomic model literally, then we should not allow for first order correlations in the increasing hazard model. Interestingly, if we re-estimate the increasing hazard model setting the 21 first order serial correlation coefficients to zero, the likelihood is still substantially better than in the partial adjustment model with unconstrained serial correlation (2,438 versus 2,405) and in the non-structural second order-autoregressions model (2,438 versus 2,434). Two nonlinear parameter (λ_2 and σ_ϵ), common across sectors, do better than 21 and 42 free sectoral serial correlation coefficients.

---

[34] And contemporaneous correlation between innovations, as in the preceding columns.

[35] The partial adjustment coefficient is equal to $1 - e^{-\lambda_0}$.

[36] The likelihood has a second maximum which is slightly better than the one reported here (likelihood improvement of 2.7) but which corresponds to an implausible partial adjustment coefficient close to one and very large unexplained serial correlation of aggregate shocks. Since the third column has a substantially better likelihood than either one of these maxima, and the underlying model may be interpreted as that of a partial adjustment with adjustment parameter equal to one and AR(2) innovations, we concentrate on the structurally meaningful estimates in the case of the first column.

[37] Since σ_ϵ is not identified in column 1, we cannot formally use the standard χ^2 test to assess the improvement in the fit as we allow for $\lambda_2 > 0$. This is the "nuisance" parameter problem discussed in, e.g., Hansen (1993). This is a second order problem in our case, however, where the likelihood differences are large.

[38] Strictly speaking, calling λ_0 the "linear" term is appropriate only in conjunction with the approximation $(e^{-x} - 1) \approx -x$ in equation (11). Since our main conclusions do not depend on this caveat, we disregard it in what follows.

Table 1: Semi-structural results

Parameters	Equipment			Structures		
	P.A.M.	Incr. Haz.	AR(2)	P.A.M.	Incr. Haz.	AR(2)
λ_0	0.582* (0.125)	0.000 (0.031)	—	0.501* (0.098)	0.000 (0.203)	—
λ_2	—	1.035 (0.318)	—	—	0.644 (0.198)	—
σ_ϵ	—	0.116 (0.036)	—	—	0.051 (0.019)	—
$\bar{a}_1$	0.19	0.04	0.75	0.07	−0.10	0.63
$\bar{a}_2$	—	—	−0.14	—	—	−0.07
Likelihood	2404.8	2444.2	2434.4	2551.7	2643.0	2588.6

Note: * stands for estimates obtained with the approximation $e^{-(x-v-\delta)} - 1 \approx v + \delta - x$ in equation (11). The autoregressive model for the sectoral residuals is: $v_{it} = const. + a_{1i}v_{it-1} + a_{2i}v_{it-2} + e_{it}$.

The results for structures are qualitatively similar to those for equipment. The first column shows a coefficient λ_0 which is significantly different from zero and implies a partial adjustment coefficient of 0.39. The mean first order serial correlation coefficient is 0.07, with a range between -0.12 and 0.27. The second column indicates the relevance of considering non-constant hazards: The likelihood function exhibits a substantial improvement when λ_2 is allowed to be positive, while the linear term, λ_0, again becomes indistinguishable from zero. The first order correlations of sectoral shocks remain on average close to zero (they range from -0.30 to 0.07). Finally, the third column shows that unconstrained (across sectors) second-order autoregressions lead to a worse fit than the simplest increasing hazard model.[39]

4.3 Results: Structural estimation

Rather than estimating the adjustment hazard directly, in this section we estimate the parameters of the adjustment costs function and obtain the hazard from the solution to the dynamic optimization problem presented in section 2. We do not estimate all the structural parameters, however. We assume an interest rate, share of capital, and markup of 6, 30 and 20 percent, respectively;[40] as well as depreciation rates for equipment and structures of 10 and 5 percent per year, respectively.

Adjustment costs are drawn from a Gamma distribution:

$$G(\omega) = \frac{1}{\phi^p \Gamma(p)} \int_0^{\omega} \eta^{p-1} e^{-\eta/\phi}\, d\eta,$$

which has mean $\mu_\omega = p\phi$ and a coefficient of variation $\text{cv}_\omega = 1/\sqrt{p}$. We estimate μ_ω, cv_ω, and σ_ϵ. As before, we allow for a general variance-covariance matrix for sectoral innovations, and sectoral first order serial correlation coefficients. The initial cross section distribution is taken to be the ergodic one but, as before, we exclude the first three observations of each series in the likelihood in order to reduce any systematic effect of this approximation.[41]

The results are reported in Table 2. For equipment, adjustment costs are on average

---

[39] As before, even if the 21 first order serial correlation coefficients are set to zero in the increasing hazard model, the likelihood is substantially larger than in the partial adjustment model with unconstrained serial correlation (2,634 versus 2,552), and in the non-structural second-order autoregressions model (2,634 versus 2,589). Once again, two nonlinear parameters, common across sectors, do better than 21 and 42 free sectoral serial correlation coefficients.

[40] These parameters imply $\beta = 0.6$.

[41] We approximate the ergodic density by the density that results after 30 iterations, beginning with a point mass at zero.

Table 2: Structural results

Parameters	Equipment	Structures
μ_ω	0.165	0.587
	(0.016)	(0.080)
cv_ω	0.342	0.143
	(0.082)	(0.027)
$\bar{\sigma}_\epsilon$	0.108	0.074
	(0.060)	(0.013)
$\bar{\rho}_1$	0.131	−0.005
σ_{ρ_1}	(0.118)	(0.095)
Likelihood	2446.6	2649.2
λ^-	3.68	4.525
λ^+	12.64	20.96
x^-	−0.472	−0.850
x^+	0.398	0.693

16.5 percent of annual revenues (net of labor payments) and their coefficient of variation is 0.34, which implies a standard deviation of 5.6 percent. Average serial correlation is small and the likelihood is slightly larger than that obtained with the semi-structural inverted-normal hazard of the previous section.[42] Indeed, our semi-structural representation of the structural hazard takes the form given in equation (25). The rows of the lower panel in Table 2 report the implicit value of the parameters of this representation.

For structures, adjustment costs are on average 58.7 percent of annual revenues (net of labor payments) and their coefficient of variation is 0.143, which implies a standard deviation of 8.4 percent. Average serial correlation is negligible, and the likelihood is larger than that obtained with the semi-structural inverted normal distribution of the previous section.[43] As with equipment, estimating the semi-structural hazard in (25) without the constraints imposed by the microeconomic dynamic optimization problem does not raise the likelihood significantly.

---

[42] For consistency with the microeconomic model, it is important that the serial correlation coefficients be reasonably close to zero.

Figure 6 shows the implied adjustment hazards (panel a) for equipment (solid line) and structures (dashed line), and the corresponding estimated distribution of adjustment costs (panel b). The hazards, which are depicted in the space of percentage deviations from the target levels, imply a substantial range where no adjustment occurs; they are sharply increasing thereafter. Clearly, there is more inaction for structures than for equipment. This difference is larger than can be inferred directly from the figure, since the depreciation and idiosyncratic uncertainty are smaller for structures than for equipment, thereby making large disequilibria less likely for structures than for equipment (for a given hazard). Our model's explanation for this difference in inaction is found in the bottom panel, where the distributions of adjustment costs are illustrated: adjustment costs are on average much larger for structures than for equipment investment.

4.4 What is special of increasing hazard models?

Perhaps the main distinctive feature of the model we have estimated, compared with its linear counterparts, is that not only average investment by those that are investing but also the number of firms that choose to invest at any point in time fluctuates over the business cycle. That this is realistic is confirmed by the evidence in Doms and Dunne (1993), who use the Longitudinal Research Datafile to study the investment behavior of 12,000 continuing U.S. manufacturing establishments for the seventeen year period from 1972-1988. Among many interesting facts, they show that the number of plants going through their primary investment spikes (i.e., the single year with the largest investment for the establishment), rather than the average size of these spikes, tracks closely aggregate manufacturing investment over time. In terms of our model, this flexibility in the number of firms investing implies that, contrary to the case of the linear model, the extent of the response of aggregate investment to aggregate shocks fluctuates over the business cycle.

Figure 7 depicts the paths of the derivatives of aggregate investment with respect to aggregate shocks for equipment and structures, evaluated at v_t. It is apparent that this "index

---

[43] Therefore substantially larger than for the linear partial adjustment and second-oder autoregression models.

Figure 6a: Adjustment Hazards

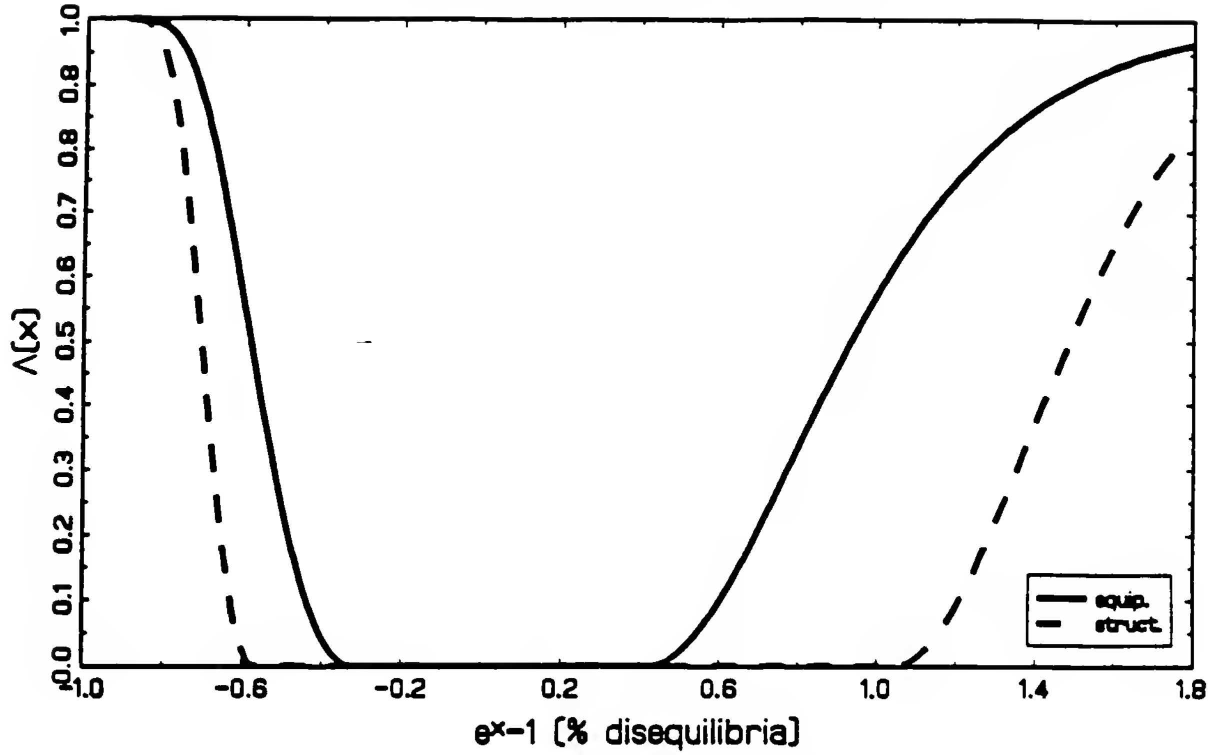

Figure 6b: Distributions of adjustment costs

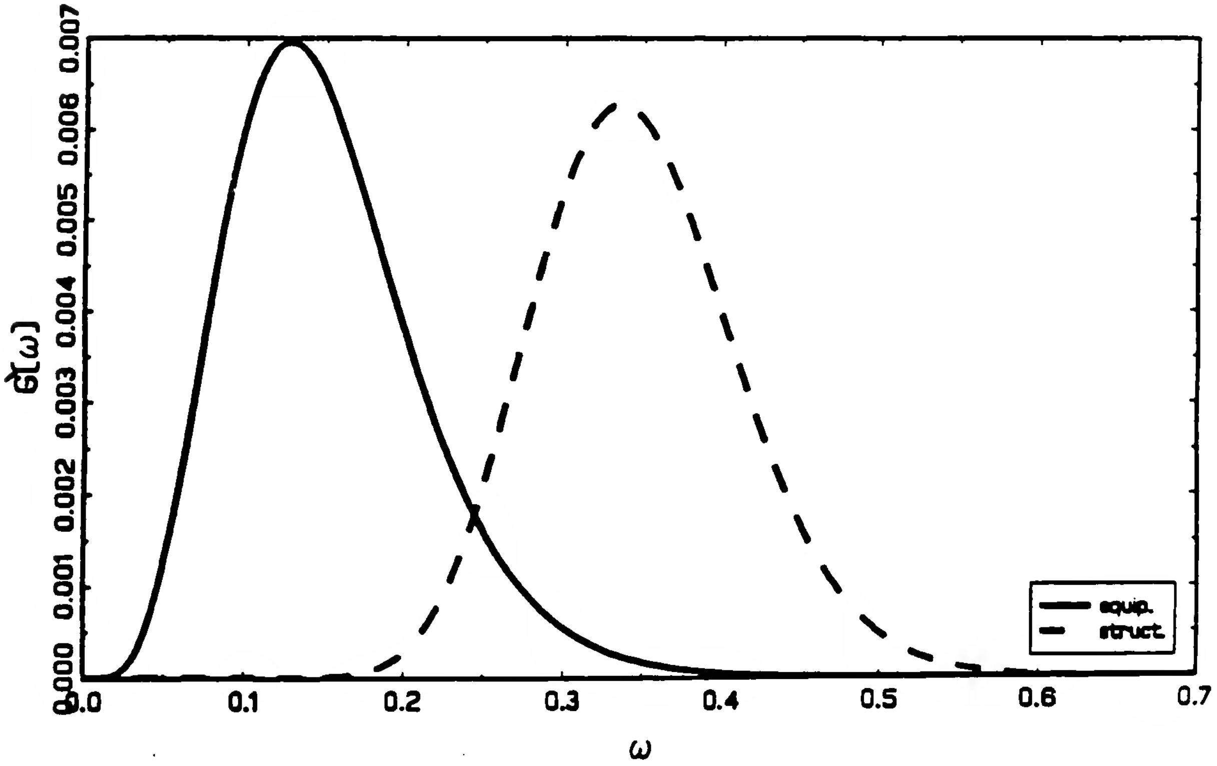

Figure 7a: Equipment-index of responsiveness

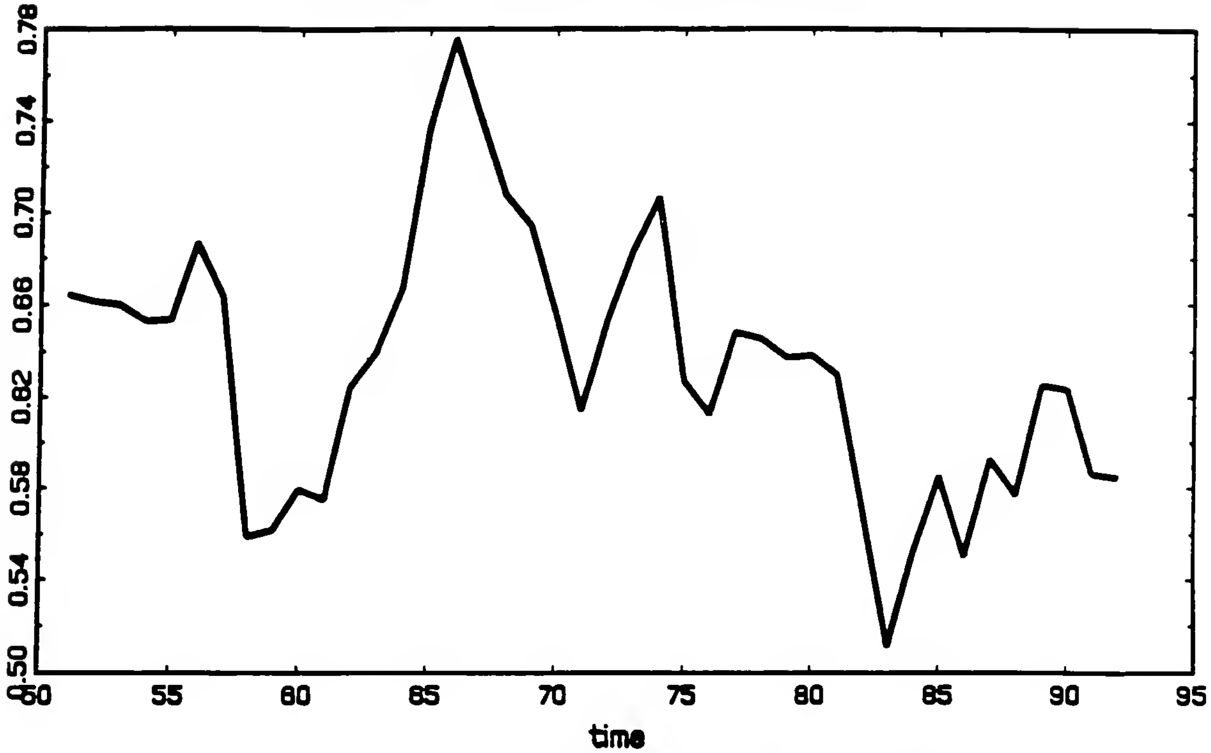

Figure 7b: Structures-index of responsiveness

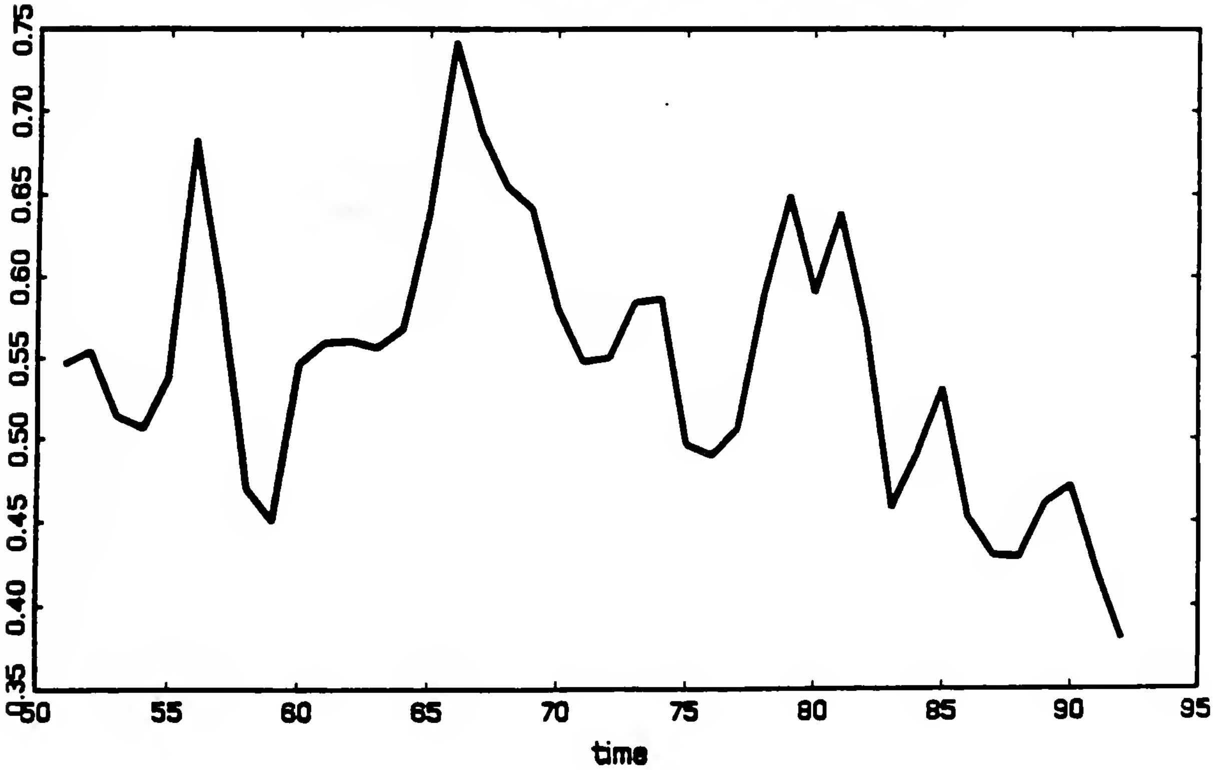

Table 3: Skewness and kurtosis for innovations*

Model	Equipment		Structures	
	Skewness	Kurtosis	Skewness	Kurtosis
Part. Adj.	0.36 (< 0.01)	1.18 (< 0.01)	0.74 (< 0.01)	1.75 (< 0.01)
AR(2)	0.35 (< 0.01)	0.85 (< 0.01)	0.81 (< 0.01)	1.72 (< 0.01)
Structural	−0.10 (> 0.20)	0.65 (< 0.01)	0.07 (> 0.20)	0.33 (> 0.07)

* Shown in parenthesis: p-values obtained via bootstrap.

of responsiveness" fluctuates widely over the sample. Moreover, it is strongly procyclical: Its correlation with aggregate shocks is 0.79 for equipment and 0.72 for structures.

Our model implies a complex non-linear process for investment. To a first approximation, however, the contemporaneous response of investment to aggregate shocks can be characterized in simpler terms. The procyclicality of the index of responsiveness described above implies that for any given sequence of shocks, the increasing hazard model has sharper (non-linear) cyclical features. In particular, it will generate fatter tails, especially so during expansions (the asymmetry is due to the strong drift induced by depreciation). Figure 8 shows the mean difference (normalized by average I/K) between actual investment and the predictions of the linear model; the latter supplied with the sequence of shocks inferred from the non-linear model. The largest absolute values of this series occurs during periods of large departure between investment and its mean.[44]

. An alternative way to reach a similar conclusion with respect to the role of nonlinearities, is to look directly at the departures of I/K and shocks generated by each model from normality. Figure 9 shows the histogram of standardized sectoral I/K.[45] It is apparent that I/K is not normal; its skewness and (excess) kurtosis coefficients are 0.61 and 0.74 for equipment and 0.76 and 0.87 for structures. Obviously, linear models with normal

---

[44] If we normalize by actual rather than average I/K, the difference during recessions is accentuated.

[45] Standardized within each sector.

Figure 8a: Equipment-mean-difference

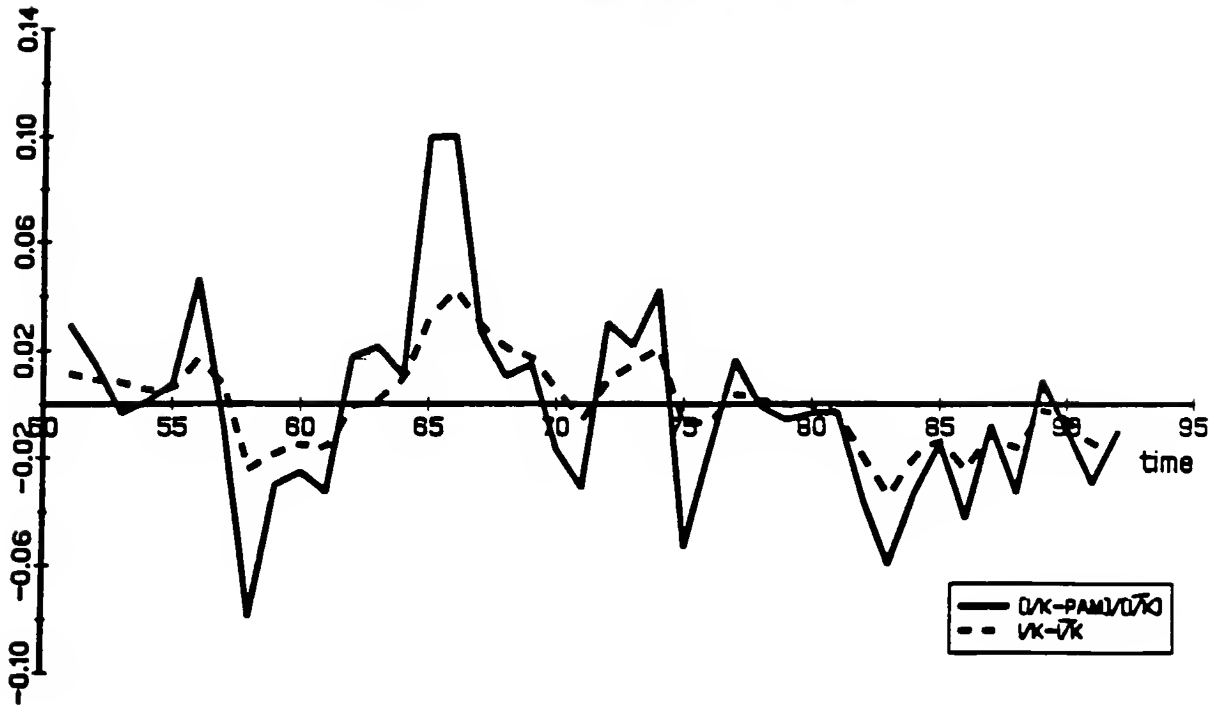

Figure 8b: Structures-mean-difference

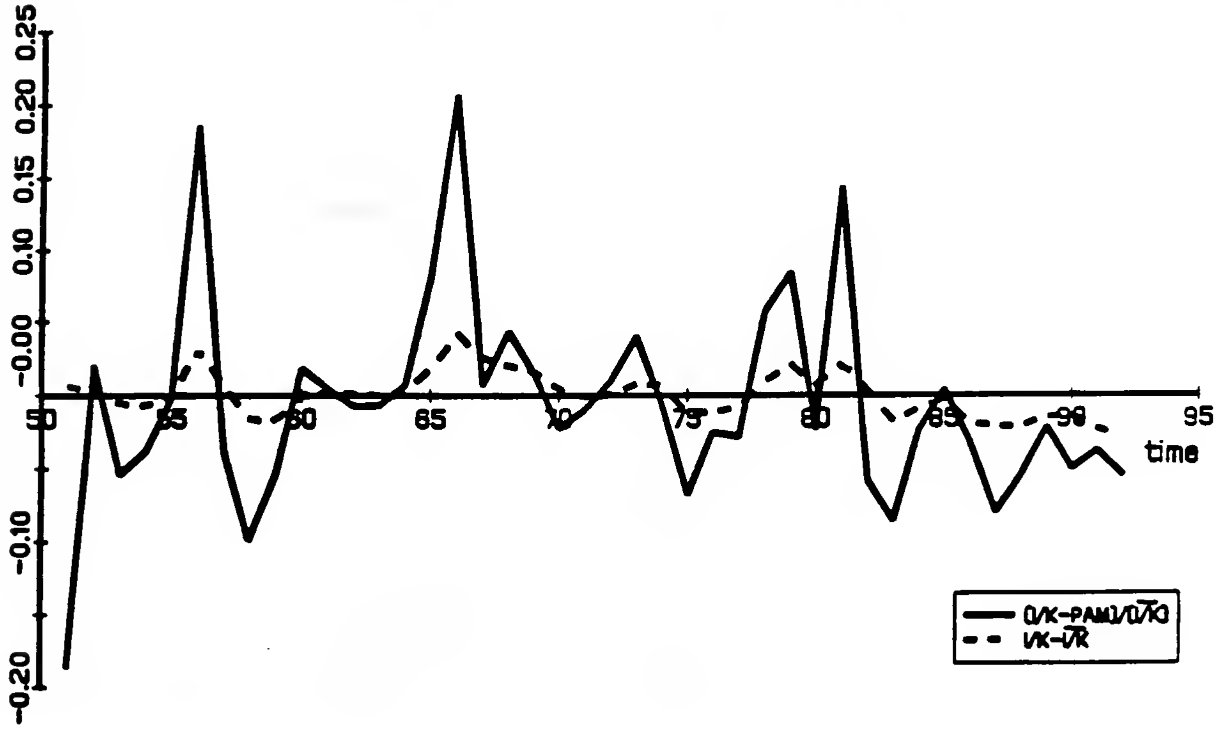

Figure 9a: Histogram of Equipment I/K (standardized)

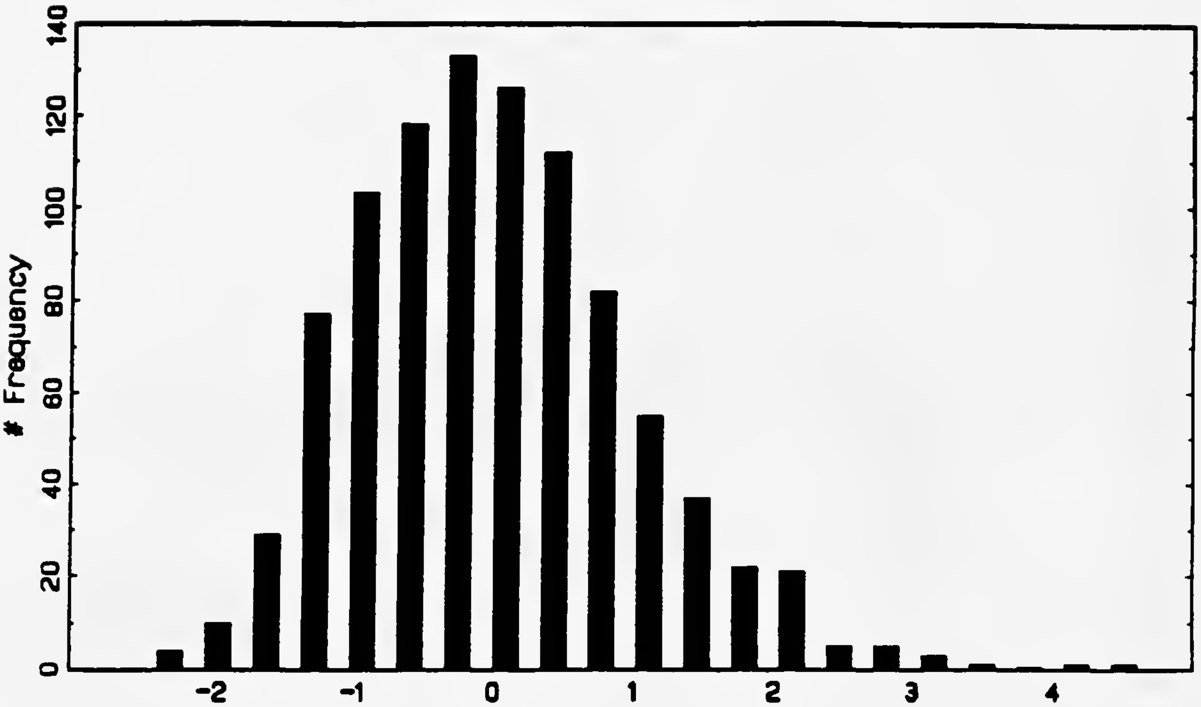

Figure 9b: Histogram of Structures I/K (standardized)

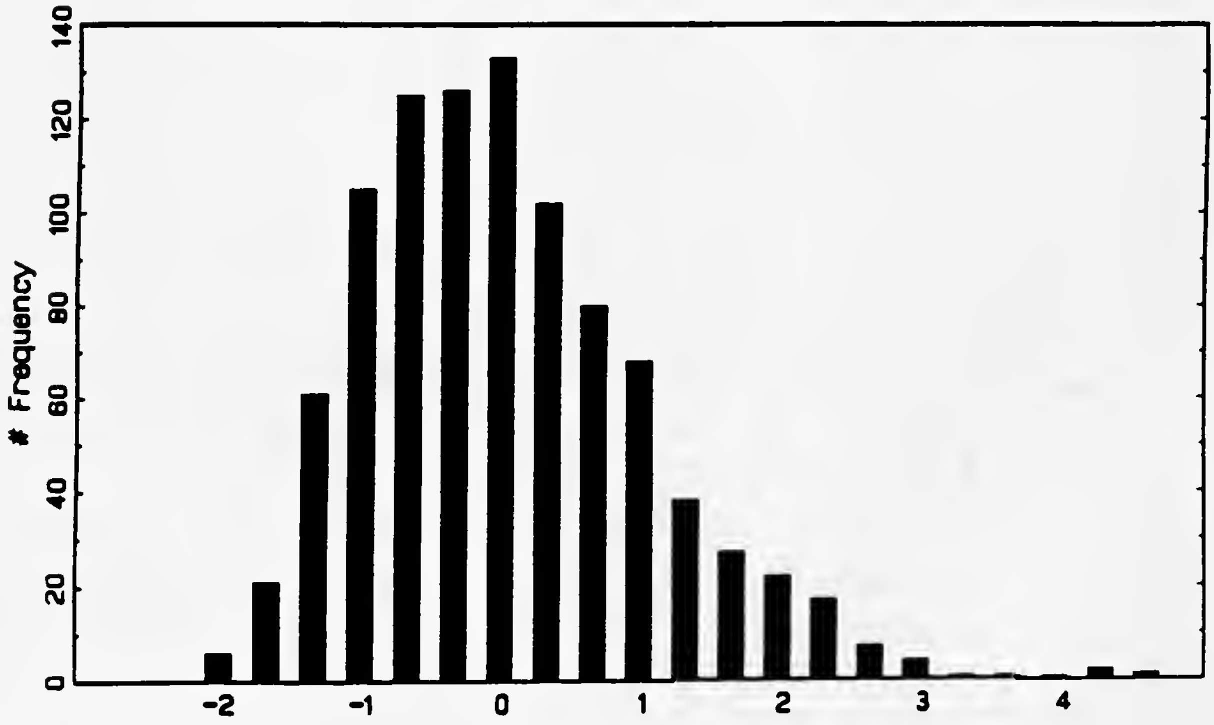

errors cannot account for these departures. The innovations generated by the best partial adjustment model and best second-order autoregression models also depart from normal, as can be seen in table 3: Their skewness and kurtosis coefficients are 0.36 and 1.18 for equipment and 0.74 and 1.75 for structures in the partial adjustment case, and 0.35 and 0.85 for equipment, and 0.81 and 1.72 for structures in the AR(2) case. All these numbers are significantly different from zero (the normal case) at the 0.01 level. The last two rows show that the increasing hazard model generates innovations that are closer to normal than its linear counterparts. The estimated skewness and kurtosis coefficients are considerably smaller, and three out of four of the coefficients do not depart significantly (at the 0.05 level) from their values under the normality assumption. The non-linear model does not need to introduce nearly as much skewness and kurtosis in aggregate shocks to account for investment behavior.

5 SMALL CHANGES

The models estimated in the preceeding section (see Figure 6) imply that firms *always* adjust their capital by large amounts (at least 30%). Strictly interpreted, this implication is unrealistic at the microeconomic level, for in addition to large projects, plants often experience small adjustments.[46]

In this section we extend the model developed earlier and incorporate the possibility of small adjustments. We consider two kinds of capital, one which is costly to adjust and another which can be adjusted costlessly. We estimate the extended model and conclude that, even though now most adjustments at the firm level are small, the occasional lumpy nature of adjustments continues playing a central role for understanding aggregate dynamics. We conclude that it is not excluding the possibility of small adjustments, but allowing for lumpy investment, that drives the results in the preceeding sections.

5.1 Model Extension

The model of Section 2 is extended to allow for two kinds of capital, one with adjustment costs ("fixed") and the other without ("flexible").[47] The former is denoted by K and the latter by k. Firms' technology is Cobb-Douglas with constant returns, now with three inputs instead of two. The elasticity of production with respect to both kinds of capital is

---

[46] See Doms and Dunne (1993).

[47] Fixity and flexibility only refer to the presence or absence of adjustment costs.

$\alpha\tau$ and $\alpha(1-\tau)$, respectively, with $0 \leq \tau \leq 1$. Thus, the model of section 2 corresponds to the particular case where $\tau = 1$.

It is easy to show that a firm with a stock of fixed capital equal to K and a deviation from its frictionless optimal stock of fixed capital of z maximizes its profits by choosing:

$$k = \mathrm{e}^{-(1-\beta)z} K, \tag{27}$$

with

$$\beta = \frac{\tau\alpha(\eta-1)}{1+\tau\alpha(\eta-1)}. \tag{28}$$

Thus, when the firm's desired stock of fixed capital increases (i.e., $\Delta z < 0$) and the current adjustment cost is too high to make adjustment profitable, the firm will adjust its flexible capital, k, even if it does not adjust its fixed capital. In every period the firm adjusts its flexible capital in the direction determined by the current change in desired frictionless capital – the magnitude of this change is larger when fixed capital is also adjusted.

Expressing k as a function of K, and using the expression for β in (28), equation (1) and the arguments that follow in section 2.2 apply directly to the more general case considered here. Thus the shape of the firm's optimal policy is the same as before: conditional on the current draw of adjustment cost factor, the firm follows an (S, s) rule for adjusting its fixed capital. Its flexible capital adjusts according to (27).

5.2 Sectoral Investment

Even though calculations are somewhat more cumbersome, aggregate equations are obtained following analogous steps to those of section 2.3. The evolution of the cross-section of firms' disequilibria is identical to what we saw in that section. The extension of the model only requires modifying the expressions relating how a given cross-section of disequilibria translates into aggregate investment. We sketch the steps required to do this next.

Let:

$$\nu = \left[\frac{1-\tau}{\tau}\right] \mathrm{e}^{-(1-\beta)c},$$

with c denoting the disequilibrium immediately after firms adjust their stock of fixed capital. Let I_t^K and I_t^k denote aggregate investment during period t in fixed and flexible capital, respectively, and let I_t^T and K_t^T denote total investment and total capital stock, respectively.

Then:

$$
\begin{aligned}
\frac{I_t^K}{K_t^T} &= \left(\frac{1}{1+J_{t-1}}\right)\int(\mathrm{e}^{-x}-1)\Lambda(x)\tilde{f}(x,t)dx, \\
\frac{I_t^k}{K_t^T} &= \left(\frac{1}{1+J_{t-1}}\right)\left[(J_t-J_{t-1})+J_t\int(\mathrm{e}^{-x}-1)\Lambda(x)\tilde{f}(x,t)dx\right],
\end{aligned}
$$

where:

$$J_t = \nu\int \mathrm{e}^{-(1-\beta)x}f(x,t)dx, \tag{29}$$

with $f(x,t)$ and $\tilde{f}(x,t)$ defined in section 2.3.2.

Adding both expressions above we obtain $y_t = I_t^T/K_t^T$. With this expression at hand, the modifications needed in section 3 when deriving the non-linear equations that must be solved to find the aggregate (sectoral) shocks, v_t, are straightforward.

5.3 Results

We estimate structural models for equipment and structures separately, as in section 4.3.[48,49] Table 4 presents the results.[50]

The parameter measuring the share of both kinds of capital, τ, is significantly different from 1 (the value implicit in section 4), but large enough so the fixed capital continues playing a key role, especially so for structures. The remaining structural parameters (mean and coefficient of variation of the Gamma distribution characterizing adjustment costs, and the standard deviation of idiosyncratic shocks) take similar values to those obtained earlier. The first order correlations are larger than in the preceeding section, specially for equipments. This may be due to assuming that flexible capital involves no adjustment costs; if the corresponding adjustment costs are smaller than for fixed capital, but non-negligible, then a first order correlation term may be expected.

More than the structural interpretation, what is important to stress here is that even though over 90 percent of yearly investment is smaller than 20 percent, the nonlinearities at the microeconomic and aggregate level are fully accounted for by the large adjustments.

---

[48] Thus in each case there are two kinds of capital: with and without adjustment costs.

[49] The interest rate, (total) share of capital (α), and markup are fixed at the same values as before.

[50] We include an additive constant (common across sectors) which partly captures investment that is totally unrelated to our model; the estimated values are 0.045 for equipment and 0.020 for structure. As mentioned earlier, given the non-linear nature of the model, it is computationally infeasible to combine unobserved, stochastic aggregate shocks with measurement error.

Table 4: Extended model: results

Parameters	Equipment	Structures
tau	0.563	0.829
	(0.103)	(0.098)
μ_ω	0.184	0.522
	(0.011)	(0.099)
cv_ω	0.235	0.172
	(0.029)	(0.101)
σ_ϵ	0.038	0.055
	(0.006)	(0.010)
$\bar{\rho}_1$	0.283	0.069
σ_{ρ_1}	(0.111)	(0.087)
Likelihood	2451.3	2660.4

And these are the ones that explain the improved aggregate fit of the models presented in this paper over standard linear ones.

6 FINAL REMARKS

In this paper we derived and estimated a model of sectoral investment that builds from the realistic observation that lumpy adjustments play an important role in firms' investment behavior, but that allows for the empirically appealing feature that adjustments do not need to be of the same size across adjusting firms and for a firm over time.

Using a non-linear time series procedure, we estimated the distributions of fixed adjustment costs faced by firms that maximizes the likelihood of *aggregate* (sectoral) data. Restricting these distributions to the gamma-family, we found that their means and standard deviations are 16.5 and 5.6 percent of a years' revenue (net of labor payments) for equipment, and 58.7 and 8.4 percent for structures. More importantly, the adjustment hazards implied by these findings are clearly non-linear: they leave a significative range of complete inaction, and increase sharply thereafter.

At the aggregate level, the estimated hazards imply brisk cyclical features. These non-linearities noticeably improve the aggregate performance of investment equations.

The empirical methodology developed in this paper should serve as a useful complement to microeconomic studies of investment, as well as in many other applications where intermittent microeconomic adjustment is suspected.

References

[1] Bar-Ilan, Avner, and Alan S. Blinder. "Consumer Durables: Evidence on the Optimality of Doing Nothing." *Journal of Money, Credit and Banking*, **24**, 253–272, 1992.

[2] Bernanke, Benjamin S., "The Determinants of Investment: Another Look," *American Economic Review* **73**, 71–75, 1983.

[3] Bertola, Giuseppe, and Ricardo J. Caballero, "Kinked Adjustment Costs and Aggregate Dynamics," in *NBER Macroeconomics Annual 1990* O.J. Blanchard and S. Fischer, eds. Cambridge, MA: The MIT Press, 1990.

[4] Bertola, Giuseppe, and Ricardo J. Caballero "Irreversibility and Aggregate Investment," *Review of Economic Studies 61*, 1994, p. 223-246.

[5] Blanchard, Olivier J., "Investment, Output, and the Cost of Capital: A Comment," *Brookings Papers of Economic Activity*, 153–158, 1986.

[6] Blinder, Alan S., "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity* II, 443–505, 1981.

[7] Caballero, Ricardo J. "Durable Goods: An Explanation for their Slow Adjustment." *Journal of Political Economy* **101**, 351–384, 1993.

[8] Caballero, Ricardo J., and Eduardo M.R.A. Engel. "Beyond the Partial Adjustment Model." *American Economic Review, Papers and Proceedings* 82, 360–364, 1992a.

[9] Caballero, Ricardo J., and Eduardo M.R.A. Engel, "Price Rigidities, Asymmetries, and Output Fluctuations," NBER WP 4091, 1992b.

[10] Caballero, Ricardo J., and Eduardo M.R.A. Engel, "Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy," *Review of Economic Studies*, **60**, 95–119, 1993a.

[11] Caballero, Ricardo J., and Eduardo M.R.A. Engel, "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, **108-2**, 697–711, 1993b.

[12] Caballero, Ricardo J., and Eduardo M.R.A. Engel, "Microeconomic Rigidities and Aggregate Price Dynamics," *European Economic Review*, **37 (4)**, 697–711, June 1993c.

[13] Caballero, Ricardo J., Eduardo M.R.A. Engel and John Haltiwanger, "Aggregate Employment Dynamics: building from microeconomic evidence," mimeo MIT, 1993.

[14] Chirinko, Robert S., "Business Fixed Investment Spending: A Critical Survey of Modelling Strategies, Empirical Results, and Policy Implications," *Journal of Economic Literature*, **31** (**4**), 1875–1911, 1993.

[15] Clark, Peter K., "Tax Incentives and Equipment Investment," *Brookings Papers on Economic Activity 1*, 317–339, 1993.

[16] Cooper, Russell and John Haltiwanger, "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review 83-3*, 360–382, 1993.

[17] Cummins, Jason G. Kevin H. Hassett, and R. Glenn Hubbard, "A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments," Mimeo Columbia, February 1994.

[18] Dixit, Avinash, and Robert S. Pindyck, *Investment Under Uncertainty*, Princeton University Press, 1994.

[19] Doms, M. and T. Dunne, "An Investigation into Capital and Labor Adjustment at the Plant Level," mimeo, Center for Economic Studies, Census Bureau, 1993.

[20] berly, J.C., "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," *Journal of Political Economy 102-3*, June 1994, p. 403-436.

[21] Grossman, Sanford J., and Guy Laroque, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, **58** (**1**), 25–51, 1990.

[22] Hall, Robert E. and Dale Jorgenson, "Tax Policy and Investment Behavior," *American Economic Review*, **57**, 391–414, 1967.

[23] Hansen, Bruce E., "Inference When a Nuisance Parameter is Not Identified Under the Null Hypothesis," mimeo, Rochester, June 1993.

[24] Harrison, J. Michael, Thomas M. Sellke and Allison J. Taylor, "Impulse Control of Brownian Motion," *Mathematics of Operations Research*, **8**, 454–466, 1983.

[25] Lam, Sam-Pok, "Permanent Income, Liquidity, and Adjustment of Automobile Stocks: Evidence from Panel Data," *Quarterly Journal of Economics*,**106**, 203-230, 1991.

[26] Rothschild, Michael, "On the Costs of Adjustment," *Quarterly Journal of Economics*, **75**, 605–622, 1971.

[27] Shapiro, Matthew D., "Investment, Output, and the Cost of Capital," *Brookings Papers of Economic Activity*, **1986-1**, 111–152, 1986.

MIT LIBRARIES
3 9080 00832844 2

Date Due

AUG 3 1 1996

FEB 2 8 1999

JWA

MAR 0 2 1999

Lib-26-67